Discretized Multinomial Distributions and Nash Equilibria in Anonymous Games

Constantinos Daskalakis* Christos H. Papadimitriou†
University of California, Berkeley
Computer Science
{costis, christos}@cs.berkeley.edu

Abstract

We show that there is a polynomial-time approximation scheme for computing Nash equilibria in anonymous games with any fixed number of strategies (a very broad and important class of games), extending the two-strategy result of [16]. The approximation guarantee follows from a probabilistic result of more general interest: The distribution of the sum of n independent unit vectors with values ranging over $\{e_1, \ldots, e_k\}$, where e_i is the unit vector along dimension i of the k-dimensional Euclidean space, can be approximated by the distribution of the sum of another set of independent unit vectors whose probabilities of obtaining each value are multiples of $\frac{1}{z}$ for some integer z, and so that the variational distance of the two distributions is at most ϵ, where ϵ is bounded by an inverse polynomial in z and a function of k, but with no dependence on n. Our probabilistic result specifies the construction of a surprisingly sparse ϵ-cover — under the total variation distance — of the set of distributions of sums of independent unit vectors, which is of interest on its own right.

1 Introduction

The recent results implying that the Nash equilibrium is an intractable problem [19], even in the two-player case [11], have directed the interest of researchers towards algorithms or complexity results for special cases [25, 34, 1, 28, 18] and approximation algorithms [32, 31, 20, 24, 21, 10, 39, 18], and the following has emerged as the main open question in the area of equilibrium computation: **Is there a PTAS for the Nash equilibrium?**[1]

In this paper we make progress on this problem, focusing on a very broad and common class of games called *anonymous games* [7, 8]. A game is anonymous if the utility of each player depends not on exactly which other player chooses which strategy; instead, it only depends on the *number* of other players that play each strategy (that is, it is a symmetric function of the strategies played by other players). Anonymous games are a much more general class than the symmetric games (known to be solvable in polynomial time when the number of strategies is fixed [35]), in which all players are identical. Many problems of interest in computational game theory, such as congestion games, participation games, voting games, and certain markets and auctions, are anonymous. Anonymous games have also been used for modeling certain social phenomena [8]. Since in anonymous games a player's utility depends on the *partition* of the remaining players into strategies, such games are a rare case of multiplayer games that have a polynomially succinct representation — as long as the number of strategies is fixed. *Our main result is a PTAS for such games.* (However, it should be noted that it is not known whether this special case of the Nash equilibrium problem is PPAD-complete, and so even an exact algorithm may be possible.)

Our PTAS extends to several generalizations of anonymous games, for example the case in which there are a few *types* of players, and the utilities depend on how many players *of each type* play each strategy; and to the case in which we have *extended families* (disjoint graphical games of constant degree and with up to logarithmically many players, each with a utility depending in arbitrary, possibly non-anonymous, ways on their neighbors, in addition to their anonymous, possibly typed, interest on everybody else). Essentially any further extension leads to intractability.

Algorithmic Game Theory aspires to understand the Internet and the markets it encompasses and creates, and therefore it should focus on *multi-player* games. We believe that our PTAS is a positive algorithmic result spanning a vast expanse in this space. However, because of the tremendous analytical difficulties detailed below, our algorithm is

*Supported by a Microsoft Research Fellowship.
†The authors were supported through NSF grant CCF - 0635319, a gift from Yahoo! Research, and a MICRO grant.
[1] It is shown in [12] that an FPTAS is no more likely than an exact solution.

not practical (as we shall see, the number of strategies appears, exponentially, in the exponent of the running time). It could be, of course, the precursor of more practical algorithms (in fact, such an algorithm for the two-strategy case has been recently proposed [15]). But, more importantly, our algorithm should be seen as compelling computational evidence that there are very extensive and important classes of common games which are free of the negative implications of the complexity result in [19].

The basic idea of our algorithm is extremely simple and intuitive (and in fact it had been noted in the past [29]): Since we are looking for mixed strategies (probability distributions, one for each player, on the set of strategies) that are in equilibrium, we restrict our search to probability distributions assigning to the strategies probabilities that are multiples of a fixed fraction, call it $\frac{1}{z}$, where z is a large enough natural number. We call this process *discretization*. We can then consider each discrete probability distribution as a separate strategy and look for (approximate) *pure* equilibria in the resulting game (the utilities of the new game can be computed via dynamic programming). The challenge is to prove that any mixed Nash equilibrium of the original game has to be close to some approximate pure Nash equilibrium of the resulting game. For general games this is not very hard to see (even though it had apparently escaped the attention of the researchers who first suggested the discretization method [29]), and this observation yields a $N^{O\left(\log \frac{\log N}{\epsilon}\right)}$ *quasi*-PTAS for computing Nash equilibria in games in which all players have a fixed number of strategies, where N is the size of the input (Theorem 4.2; note that this complements the $N^{O\left(\log N/\epsilon^2\right)}$ quasi-PTAS of [32] for games with a fixed number of *players*). We also point out that the discretization method gives the first algorithm for tree-like graphical games with a fixed number of strategies (for trees, an initial attempt by [30] in the two-strategy case was found to have flaws in [23], while in the latter paper a polynomial-time algorithm for graphical games with two strategies on *paths* and *cycles* was developed). Our algorithm applies to all graphical games with a fixed number of strategies whose graph is of bounded degree and logarithmically bounded treewidth.

The discretization method requires polynomial time in the case of anonymous games, because in this case the search space is no longer the set of all n-tuples of discrete distributions, where n is the number of players (this is exponential in n); instead, via dynamic programming (see the proof of Theorem 2.2), it can be reduced to the set of all the ordered partitions of n into $\ell = O((z+1)^{k-1})$ parts, where ℓ is the number of discrete probability distributions defined above, which is polynomial in n, if k, the number of strategies, and z, the discretization, are fixed.

But proving in this case that the approximation is valid turns out to be a deep problem. One has to establish a probabilistic lemma stating that, given a multinomial-sum distribution (the sum of k-dimensional unit vector-valued independent but not necessarily identically distributed random variables), the probabilities can be rounded to multiples of $\frac{1}{z}$ so that the variational distance between the resulting distribution and the original one depends only on z (and in fact this dependence is inversely polynomial), and on the dimension k (in an arbitrary way; the bound we can prove is exponential, and we suspect it is necessary). This probabilistic lemma for the case of two strategies (i.e., for binomial-sum distributions) was proved in [16] by clustering the variables into three classes, depending on how large their expectation is, and then using results from the probability literature [4, 5, 37] to approximate each component binomial-sum distribution (both the original and the rounded one) by Poisson or shifted Poisson distributions (depending on the cluster), and finally rounding the probabilities so that the approximations are close.

In the multinomial case, however, no useful approximations are known; see, e.g., [2] for some obstacles in extending the existent methods to the multinomial case. Another reason that makes the binomial case easy is that it is essentially one-dimensional: in the multinomial case on the other hand, watching the balls in one bin, so to speak, provides small information about the distribution of the remaining balls in the other bins, because the random vectors are not identically distributed. Our proof is very involved and indirect, resorting to an alternative sampling of each random vector by funneling a ball down a probabilistic decision tree with $k-1$ leaves (k is the dimension, or number of strategies), ending up eventually with a binary choice at the leaves. This choice can now be discretized similarly to the binomial case — albeit with much more effort. The decision tree topologies become the clusters for the approximation, and their number (exponential in k) appears in the variation distance via a union bound, and, hence, in the exponent of the running time. We believe that this probabilistic lemma (Theorem 2.1), and its proof, represent an advance of some substance in the state of the art in this area of applied probability.

Our result can be interpreted as constructing a surprisingly sparse cover of the set of multinomial-sum distributions under the total variation distance. Covers of metric spaces have been considered in the literature of approximation algorithms, but we know of no non-trivial result working for the total variation distance or producing a cover of the required sparsity to achieve a polynomial-time approximation scheme for the Nash equilibrium in anonymous games. To show the value of our result in another context, we exhibit a family of non-convex optimization problems arising in economics that can be approximated by means of our probabilistic lemma and for which no efficient algorithm was known before. An application of our

result for this family of non-convex optimization problems is a PTAS for finding threat points in repeated anonymous games. These results are discussed in Section 5.

In the balance of this section we provide the necessary definitions. In the next section we describe the basics of the main result, including the algorithm and an overview of the proof. The main part of the proof of the probabilistic lemma is in Section 3, while in Section 4 we explore the application of our method to broad generalizations of anonymous games, as well as general (non-anonymous) games and graphical games. In Section 5 we present the application of our result to certain types of non-convex optimization problems. We conclude with a discussion of problems that remain open.

1.1 Definitions and Notation

An *anonymous game* is a triple $G = (n, k, \{u_i^p\})$ where $[n] = \{1, \ldots, n\}$, $n \geq 2$, is the set of players, $[k] = \{1, \ldots, k\}$, $k \geq 2$, is the set of strategies, and u_i^p with $p \in [n]$ and $i \in [k]$ is the utility of player p when she plays strategy i, a function mapping the set of partitions $\Pi_{n-1}^k = \{(x_1, \ldots, x_k) : x_i \in \mathbb{N}_0 \text{ for all } i \in [k], \sum_{i=1}^k x_i = n-1\}$ to the interval $[0,1]$. [2] Our working assumptions are that n is large and k is fixed; notice that, in this case, anonymous games are *succinctly representable* [35], in the sense that their representation requires specifying $O(n^k)$ numbers, as opposed to the nk^n numbers required for general games (arguably, succinct games are the only multiplayer games that are computationally meaningful, see [35] for an extensive discussion of this point). The convex hull of the set Π_{n-1}^k will be denoted by $\Delta_{n-1}^k = \{(x_1, \ldots, x_k) : x_i \geq 0 \text{ for all } i \in [k], \sum_{i=1}^k x_i = n-1\}$.

A *pure strategy profile* in such a game is a mapping S from $[n]$ to $[k]$. A pure strategy profile S is an *ϵ-approximate pure Nash equilibrium*, where $\epsilon \geq 0$, if, for all $p \in [n]$, $u_{S(p)}^p(x[S,p]) + \epsilon \geq u_i^p(x[S,p])$ for all $i \in [k]$, where $x[S,p] \in \Pi_{n-1}^k$ is the partition $(x_1, \ldots, x_k)$ such that x_i is the number of players $q \in [n] - \{p\}$ with $S(q) = i$.

A *mixed strategy profile* is a set of n distributions $\{\delta_p \in \Delta^k\}_{p \in [n]}$, where by Δ^k we denote the $(k-1)$-dimensional simplex, or, equivalently, the set of distributions over $[k]$. A mixed strategy profile is an *ϵ-Nash equilibrium* if, for all $p \in [n]$ and $j, j' \in [k]$,

$$E_{\delta_1, \ldots, \delta_n} u_j^p(x) > E_{\delta_1, \ldots, \delta_n} u_{j'}^p(x) + \epsilon \Rightarrow \delta_p(j') = 0,$$

where x is drawn from Π_{n-1}^k by drawing $n-1$ random samples from $[k]$ independently according to the distributions $\delta_q, q \neq p$, and forming the induced partition.

Similarly, a mixed strategy profile is an *ϵ-approximate Nash equilibrium* if, for all $p \in [n]$ and $j \in [k]$, $E_{\delta_1, \ldots, \delta_n} u_i^p(x) + \epsilon \geq E_{\delta_1, \ldots, \delta_n} u_j^p(x)$, where i is drawn from $[k]$ according to δ_p and x is drawn from Π_{n-1}^k as above, by drawing $n-1$ random samples from $[k]$ independently according to the distributions $\delta_q, q \neq p$, and forming the induced partition.

Clearly, an ϵ-Nash equilibrium is also an ϵ-approximate Nash equilibrium, but the converse is not true in general (for an extensive discussion, see [19]). All our positive approximation results are for the stronger notion of the ϵ-Nash equilibrium.

[2] In the literature on Nash approximation, utilities are usually normalized in this way so that the approximation error is additive.

2 The Main Result

The *total variation distance* between two distributions $\mathbb{P}$ and $\mathbb{Q}$ over a finite set $\mathcal{A}$ is

$$||\mathbb{P} - \mathbb{Q}||_{TV} = \frac{1}{2} \sum_{\alpha \in \mathcal{A}} |\mathbb{P}(\alpha) - \mathbb{Q}(\alpha)|.$$

Similarly, if X and Y are two random variables ranging over a finite set, their total variation distance, denoted

$$||X - Y||_{TV},$$

is defined as the total variation distance between their distributions. The bulk of the paper is dedicated to proving the following result, generalizing the one-dimensional ($k = 2$) case established in [16].

Theorem 2.1 *Let $\{p_i \in \Delta^k\}_{i \in [n]}$, and let $\{\mathcal{X}_i \in \mathbb{R}^k\}_{i \in [n]}$ be a set of independent k-dimensional random unit vectors such that, for all $i \in [n]$, $\ell \in [k]$, $\Pr[\mathcal{X}_i = e_\ell] = p_{i,\ell}$, where $e_\ell \in \mathbb{R}^k$ is the unit vector along dimension ℓ; also, let $z > 0$ be an integer. Then there exists another set of probability vectors $\{\widehat{p}_i \in \Delta^k\}_{i \in [n]}$ such that*

1. *$|\widehat{p}_{i,\ell} - p_{i,\ell}| = O\left(\frac{1}{z}\right)$, for all $i \in [n], \ell \in [k]$;*
2. *$\widehat{p}_{i,\ell}$ is an integer multiple of $\frac{1}{2^k}\frac{1}{z}$, for all $i \in [n], \ell \in [k]$;*
3. *if $p_{i,\ell} = 0$, then $\widehat{p}_{i,\ell} = 0$, for all $i \in [n], \ell \in [k]$;*
4. *if $\{\widehat{\mathcal{X}}_i \in \mathbb{R}^k\}_{i \in [n]}$ is a set of independent random unit vectors such that $\Pr[\widehat{\mathcal{X}}_i = e_\ell] = \widehat{p}_{i,\ell}$, for all $i \in [n], \ell \in [k]$, then*

$$\left|\left|\sum_i \mathcal{X}_i - \sum_i \widehat{\mathcal{X}}_i\right|\right|_{\text{TV}} = O\left(f(k)\frac{\log z}{z^{1/5}}\right) \quad (1)$$

and, moreover, for all $j \in [n]$,

$$\left|\left|\sum_{i \neq j} \mathcal{X}_i - \sum_{i \neq j} \widehat{\mathcal{X}}_i\right|\right|_{\text{TV}} = O\left(f(k)\frac{\log z}{z^{1/5}}\right), \quad (2)$$

where $f(k)$ is an exponential function of k estimated in the proof.

In other words, there is a way to quantize any set of n independent random vectors into another set of n independent random vectors, whose probabilities of obtaining each value are integer multiples of $\epsilon \in [0,1]$, so that the total variation distance between the distribution of the sum of the vectors before and after the quantization is bounded by $O(f(k)2^{k/6}\epsilon^{1/6})$. The important, and perhaps surprising, aspect of this bound is the lack of dependence on the number n of random vectors. From this, the main result of this section follows.

Theorem 2.2 *There is a PTAS for the mixed Nash equilibrium problem for anonymous games with a constant number of strategies.*

Proof: Consider a mixed Nash equilibrium $(p_1,\ldots,p_n)$. We claim that the mixed strategy profile $(\widehat{p}_1,\ldots,\widehat{p}_n)$ specified by Theorem 2.1 constitutes a $O\left(f(k)z^{-\frac{1}{6}}\right)$-Nash equilibrium. Indeed, for every player $i \in [n]$ and every pure strategy $m \in [k]$ for that player, let us track down the change in the expected utility of the player for playing strategy m when the distribution over Π^k_{n-1} defined by the $\{p_j\}_{j\neq i}$ is replaced by the distribution defined by the $\{\widehat{p}_j\}_{j\neq i}$. It is not hard to see that the absolute change is bounded by the total variation distance between the distributions of the random vectors $\sum_{j\neq i}\mathcal{X}_j$ and $\sum_{j\neq i}\widehat{\mathcal{X}}_j$, where $\{\mathcal{X}_j\}_{j\neq i}$ are independent random vectors distributed according to the distributions $\{p_j\}_{j\neq i}$ and, similarly, $\{\widehat{\mathcal{X}}_j\}_{j\neq i}$ are independent random vectors distributed according to the distributions $\{\widehat{p}_j\}_{j\neq i}$.[3] Hence, by Theorem 2.1, the change in the utility of the player is at most $O(f(k)z^{-\frac{1}{6}})$, which implies that the $\widehat{p}_i$'s constitute an $O(f(k)z^{-\frac{1}{6}})$-Nash equilibrium of the game. If we take $z=(f(k)/\epsilon)^6$, this is a δ-Nash equilibrium, for $\delta = O(\epsilon)$.

From the previous discussion it follows that there exists a mixed strategy profile $\{\widehat{p}_i\}_i$ which is of the very special kind described by Property 2 in the statement of Theorem 2.1 and constitutes a δ-Nash equilibrium of the given game, if we choose $z=(f(k)/\epsilon)^6$. The problem is, of course, that we do not know such a mixed strategy profile and, moreover, we cannot afford to do exhaustive search over all mixed strategy profiles satisfying Property 2, since there is an exponential number of those. We do instead the following search which is guaranteed to find a δ-Nash equilibrium.

Notice that there are at most $(2^k z)^k = 2^{k^2}\,(f(k)/\epsilon)^{6k} =: K$ "quantized" mixed strategies with each probability being a multiple of $\frac{1}{2^k}\frac{1}{z}$, $z=(f(k)/\epsilon)^6$. Let $\mathcal{K}$ be the set of such quantized mixed strategies. We start our algorithm by guessing the partition of the number n of players into quantized mixed strategies; let $\theta = \{\theta_\sigma\}_{\sigma\in\mathcal{K}}$ be the partition, where θ_σ represents the number of players choosing the discretized mixed strategy $\sigma \in \mathcal{K}$. Now we only need to determine if there exists an assignment of mixed strategies to the players in $[n]$, with θ_σ of them playing mixed strategy $\sigma\in\mathcal{K}$, so that the corresponding mixed strategy profile is a δ-Nash equilibrium. To answer this question it is enough to solve the following *max-flow* problem. Let us consider the bipartite graph $([n],\mathcal{K},E)$ with edge set E defined as follows: $(i,\sigma)\in E$, for $i\in[n]$ and $\sigma\in\mathcal{K}$, if $\theta_\sigma > 0$ and σ is a δ-best response for player i, if the partition of the other players into the mixed strategies in $\mathcal{K}$ is the partition θ, with one unit subtracted from θ_σ.[4] Note that to define E expected payoff computations are required. By straightforward dynamic programming, the expected utility of player i for playing pure strategy $s\in[k]$ given the mixed strategies of the other players can be computed with $O(kn^k)$ operations on numbers with at most $b(n,z,k) := \lceil 1+n(k+\log_2 z)+\log_2(1/u_{\min})\rceil$ bits, where $u_{\min}$ is the smallest non-zero payoff value of the game.[5] To conclude the construction of the max-flow instance we add a source node u connected to all the left hand side nodes and a sink node v connected to all the right hand side nodes. We set the capacity of the edge (σ,v) equal to θ_σ, for all $\sigma\in\mathcal{K}$, and the capacity of all other edges equal to 1. If the max-flow from u to v has value n then there is a way to assign discretized mixed strategies to the players so that θ_σ of them play mixed strategy $\sigma\in\mathcal{K}$ and the resulting mixed strategy profile is a δ-Nash equilibrium (details omitted). There are at most $(n+1)^{K-1}$ possible guesses for θ; hence, the search takes overall time

$$O\left((nKk^2n^k b(n,z,k) + p(n+K+2))\cdot(n+1)^{K-1}\right),$$

where $p(n+K+2)$ is the time needed to find an integral maximum flow in a graph with $n+K+2$ nodes and edge-weights encoded with at most $\lceil \log_2 n\rceil$ bits. Hence, the overall time is

$$n^{O\left(2^{k^2}\left(\frac{f(k)}{\epsilon}\right)^{6k}\right)}\cdot\log_2(1/u_{\min}).$$

■

Remark: Theorem 2.1 can be interpreted as constructing a sparse cover of the set of distributions of sums of independent random unit vectors under the total variation distance. We know of no non-trivial results working for this distance or achieving the same sparsity.

[3] To establish this bound we use the fact that all utilities lie in $[0,1]$.

[4] For our discussion, a mixed strategy σ of player i is a δ-*best response* to a set of mixed strategies for the other players iff the expected payoff of player i for playing any pure strategy s in the support of σ is no more than δ worse than her expected payoff for playing any pure strategy s'.

[5] To compute a bound on the number of bits required for the expected utility computations, note that the expected utility is positive, cannot exceed 1, and its smallest possible non-zero value is at least $(\frac{1}{2^k}\frac{1}{z})^n u_{\min}$, since the mixed strategies of all players are from the set $\mathcal{K}$.

2.1 Discussion of Proof Techniques

Observe that, from a technical perspective, the $k > 2$ case of Theorem 2.1 is inherently different than the $k = 2$ case, which was shown in [16] (Theorem 3.1). Indeed, when $k = 2$, knowledge of the number of players who selected their first strategy determines the whole partition of the number of players into strategies; therefore, in this case the probabilistic experiment is in some sense *one-dimensional*. On the other hand, when $k > 2$, knowledge of the number of "balls in a bin", that is the number of players who selected a particular strategy, does not provide full information about the number of balls in the other bins. This complication would be quite benign if the vectors $\mathcal{X}_i$ were identically distributed, since in this case the number of balls in a bin would at least characterize precisely the probability distribution of the number of balls in the other bins (as a multinomial distribution with one bin less and the bin-probabilities appropriately renormalized). But, in our case, the vectors $\mathcal{X}_i$ are not identically distributed. Hence, already for $k = 3$ the problem is fundamentally more involved than in the $k = 2$ case.

Indeed, it turns out that obtaining the result for the $k = 2$ case is easier. Here is the intuition: If the expectation of every $\mathcal{X}_i$ at the first bin was small, their sum would be distributed like a Poisson distribution (marginally at that bin); if the expectation of every $\mathcal{X}_i$ was large, the sum would be distributed like a (discretized) Normal distribution. [6] So, to establish the result we can do the following (see [16] for details): First, we cluster the $\mathcal{X}_i$'s into those with small and those with large expectation at the first bin, and then we discretize the $\mathcal{X}_i$'s separately in the two clusters in such a way that the sum of their expectations (within each cluster) is preserved to within the discretization accuracy. To show the closeness in total variation distance between the sum of the $\mathcal{X}_i$'s before and after the discretization, we compare instead the Poisson or Normal distributions (depending on the cluster) which approximate the sum of the $\mathcal{X}_i$'s: For the "small cluster", we compare the Poisson distributions approximating the sum of the $\mathcal{X}_i$'s before and after the discretization. For the "large cluster", we compare the Normals approximating the sum of the $\mathcal{X}_i$'s before and after the discretization.

One would imagine that a similar technique, i.e., approximating by a multidimensional Poisson or Normal distribution, would work for the $k > 2$ case. Comparing a sum of multinomial random variables to a multidimensional Poisson or Normal distribution is a little harder in many dimensions (see the discussion in [2]), but almost optimal bounds *are* known for both the multidimensional Poisson [2, 38] and the multidimensional Normal [6, 26] approximations. Nevertheless, these results by themselves are not sufficient for our setting: Approximating by a multidimensional Normal performs very poorly at the coordinates where the vectors have small expectations, and approximating by a multidimensional Poisson fails at the coordinates where the vectors have large expectations. And in our case, it could very well be that the sum of the $\mathcal{X}_i$'s is distributed like a multidimensional Poisson distribution in a subset of the coordinates and like a multidimensional Normal in the complement (those coordinates where the $\mathcal{X}_i$'s have respectively small or large expectations). What we really need, instead, is a multidimensional approximation result that combines the multidimensional Poisson and Normal approximations in the same picture; and such a result is not known.

Our approach instead is very indirect. We define an alternative way of sampling the vectors $\mathcal{X}_i$ which consists of performing a random walk on a binary decision tree and performing a probabilistic choice between two strategies at the leaves of the tree (Sections 3.1 and 3.2). The random vectors are then clustered so that, within a cluster, all vectors share the same decision tree (Section 3.3), and the rounding, performed separately for every cluster, consists of discretizing the probabilities for the probabilistic experiments at the leaves of the tree (Section 3.4). The rounding is done in such a way that, if all vectors $\mathcal{X}_i$ were to end up at the same leaf after walking on the decision tree, then the one-dimensional result described above would apply for the (binary) probabilistic choice that the vectors are facing at the leaf. However, the random walks will not all end up at the same leaf with high probability. To remedy this, we define a coupling between the random walks of the original and the discretized vectors for which, in the typical case, the probabilistic experiments that the original vectors will run at every leaf of the tree are very "similar" to the experiments that the discretized vectors will run. That is, our coupling guarantees that, with high probability over the random walks, the total variation distance between the choices (as random variables) that are to be made by the original vectors at every leaf of the decision tree and the choices (again as random variables) that are to be made by the discretized vectors is very small. The coupling of the random walks is defined in Section 3.5, and a quantification of the similarity of the leaf experiments under this coupling is given in Section 3.6.

For a discussion about why naive approaches such as *rounding to the closest discrete distribution* or *randomized rounding* do not appear useful, even for the $k = 2$ case, see Section 3.1 of [16].

[6] Comparing, in terms of variational distance, a sum of independent Bernoulli random variables to a Poisson or a Normal distribution is an important problem in probability theory. The approximations we use are obtained by applications of *Stein's method* [3, 4, 37].

3 Proof of Theorem 2.1

3.1 The Trickle-down Process

Consider the mixed strategy p_i of player i. The crux of our argument is an alternative way to sample from this distribution, based on the so-called *trickle-down process*, defined next.

TDP — Trickle-Down Process
Input: (S, p), where $S = \{i_1, \ldots, i_m\} \subseteq [k]$ is a set of strategies and p a probability distribution $p(i_j) > 0 : j = 1, \ldots, m$. We assume that the elements of S are ordered $i_1, \ldots, i_m$ in such a way that (a) $p(i_2)$ is the largest of the $p(i_j)$'s and (b) for $2 \neq j < j' \neq 2$, $p(i_j) \leq p(i_{j'})$. That is, the largest probability is second, and, other than that, the probabilities are sorted in non-decreasing order (ties broken lexicographically).
if $|S| \leq 2$ stop;
else apply the *partition and double operation*:

1. let $\ell^* < m$ be the (unique) index such that
   $\sum_{\ell < \ell^*} p(i_\ell) \leq \frac{1}{2}$ and $\sum_{\ell > \ell^*} p(i_\ell) < \frac{1}{2}$;
2. Define the sets
   $S_L = \{i_\ell : \ell \leq \ell^*\}$ and $S_R = \{i_\ell : \ell \geq \ell^*\}$
3. Define the probability distribution p_L such that, for all $\ell < \ell^*$, $p_L(i_\ell) = 2p(i_\ell)$. Also, let $t := 1 - \sum_{\ell=1}^{\ell^*-1} p_L(i_\ell)$; if $t = 0$, then remove ℓ^* from S_L, otherwise set $p_L(i_{\ell^*}) = t$. Similarly, define the probability distribution p_R such that $p_R(i_\ell) = 2p(i_\ell)$, for all $\ell > \ell^*$ and $p_R(i_{\ell^*}) = 1 - \sum_{\ell^*+1}^{m} p_R(i_\ell)$. Notice that, because of the way we have ordered the strategies in S, i_{ℓ^*} is neither the first nor the last element of S in our ordering, and hence $2 \leq |S_L|, |S_R| < |S|$.
4. call **TDP**(S_L, p_L); call **TDP**(S_R, p_R);

That is, TDP splits the support of the mixed strategy of a player into a tree of finer and finer sets of strategies, with all leaves having just two strategies. At each level the two sets in which the set of strategies is split overlap in at most one strategy (whose probability mass is divided between its two copies). The two sets then have probabilities adding up to $1/2$, but then the probabilities are multiplied by 2, so that each node of the tree represents a distribution.

3.2 The Alternative Sampling of $\mathcal{X}_i$

Let p_i be the mixed strategy of player i, and $\mathcal{S}_i$ be its support. [7] The execution of **TDP**$(\mathcal{S}_i, p_i)$ defines a rooted binary tree T_i with node set V_i and set of leaves ∂T_i. Each node $v \in V_i$ is identified with a pair $(S_v, p_{i,v})$, where $S_v \subseteq [k]$ is a set of strategies and $p_{i,v}$ is a distribution over S_v. Based on this tree, we define the following alternative way to sample $\mathcal{X}_i$:

SAMPLING $\mathcal{X}_i$

1. (*Stage 1*) Perform a random walk from the root of the tree T_i to the leaves, where, at every non-leaf node, the left or right child is chosen with probability $1/2$; let $\Phi_i \in \partial T_i$ be the (random) leaf chosen by the random walk;
2. (*Stage 2*) Let (S, p) be the label assigned to the leaf Φ_i, where $S = \{\ell_1, \ell_2\}$; set $\mathcal{X}_i = e_{\ell_1}$, with probability $p(\ell_1)$, and $\mathcal{X}_i = e_{\ell_2}$, with probability $p(\ell_2)$.

The following lemma, whose straightforward proof we omit, states that this is indeed an alternative sampling of the mixed strategy of player i.

Lemma 3.1 *For all $i \in [n]$, the process* SAMPLING $\mathcal{X}_i$ *outputs $\mathcal{X}_i = e_\ell$ with probability $p_{i,\ell}$, for all $\ell \in [k]$.*

3.3 Clustering the Random Vectors

We use the process **TDP** to cluster the random vectors of the set $\{\mathcal{X}_i\}_{i \in [n]}$. We define a cell for every possible tree structure. In particular, for some $\alpha > 0$ to be determined later in the proof,

Definition 3.2 (Cell Definition) *Two vectors $\mathcal{X}_i$ and $\mathcal{X}_j$ belong to the same cell if*

- *there exists a tree isomorphism $f_{i,j} : V_i \to V_j$ between the trees T_i and T_j such that, for all $u \in V_i$, $v \in V_j$, if $f_{i,j}(u) = v$, then $S_u = S_v$, and in fact the elements of S_u and S_v are ordered the same way by $p_{i,u}$ and $p_{j,v}$.*
- *if $u \in \partial T_i$, $v = f_{i,j}(u) \in \partial T_j$, and $\ell^* \in S_u = S_v$ is the strategy with the smallest probability mass for both $p_{i,u}$ and $p_{j,v}$, then either $p_{i,u}(\ell^*), p_{j,v}(\ell^*) \leq \frac{\lfloor z^\alpha \rfloor}{z}$ or $p_{i,u}(\ell^*), p_{j,v}(\ell^*) > \frac{\lfloor z^\alpha \rfloor}{z}$; the leaf is called* Type A leaf *in the first case,* Type B leaf *in the second case.*

It is easy to see that the total number of cells is bounded by a function of k only, estimated in the following claim; the proof of the claim is postponed to Appendix A.

Claim 3.3 *Any tree resulting from TDP has at most $k - 1$ leaves, and the total number of cells is bounded by $g(k) = k^{k^2} 2^{k-1} 2^k k!$.*

[7] In this section and the following two sections we assume that $|\mathcal{S}_i| > 1$; if not, we set $\widehat{p}_i = p_i$, and all claims we make in Sections 3.5 and 3.6 are trivially satisfied.

3.4 Discretization within a Cell

Recall that our goal is to "discretize" the probabilities in the distribution of the $\mathcal{X}_i$'s. We will do this separately in every cell of our clustering. In particular, supposing that $\{\mathcal{X}_i\}_{i\in\mathcal{I}}$ is the set of vectors falling in a particular cell, for some index set $\mathcal{I}$, we will define a set of "discretized" vectors $\{\widehat{\mathcal{X}}_i\}_{i\in\mathcal{I}}$ in such a way that, for $h(k) = k2^k$, and for all $j \in \mathcal{I}$,

$$\left|\left|\sum_{i\in\mathcal{I}}\mathcal{X}_i - \sum_{i\in\mathcal{I}}\widehat{\mathcal{X}}_i\right|\right|_{TV} = O(h(k)\log z \cdot z^{-1/5}); \quad (3)$$

$$\left|\left|\sum_{i\in\mathcal{I}\setminus\{j\}}\mathcal{X}_i - \sum_{i\in\mathcal{I}\setminus\{j\}}\widehat{\mathcal{X}}_i\right|\right|_{TV} = O(h(k)\log z \cdot z^{-1/5}). \quad (4)$$

We establish these bounds in Section 3.5. Using the bound on the number of cells in Claim 3.3, an easy application of the coupling lemma implies the bounds shown in (1) and (2) for $f(k) := h(k)\cdot g(k)$, thus concluding the proof of Theorem 2.1.

We shall henceforth concentrate on a particular cell containing the vectors $\{\mathcal{X}_i\}_{i\in\mathcal{I}}$, for some $\mathcal{I} \subseteq [n]$. Since the trees $\{T_i\}_{i\in\mathcal{I}}$ are isomorphic, for notational convenience we shall denote all those trees by T. To define the vectors $\{\widehat{\mathcal{X}}_i\}_{i\in\mathcal{I}}$ we must provide, for all $i \in \mathcal{I}$, a distribution $\widehat{p}_i : [k] \to [0,1]$ such that $\Pr[\widehat{\mathcal{X}}_i = e_\ell] = \widehat{p}_i(\ell)$, for all $\ell \in [k]$. To do this, we assign to all $\{\widehat{\mathcal{X}}_i\}_{i\in\mathcal{I}}$ the tree T and then, for every leaf $v \in \partial T$ and $i \in \mathcal{I}$, define a distribution $\widehat{p}_{i,v}$ over the two-element ordered set S_v, by the ROUNDING process below. Then the distribution $\widehat{p}_i$ is implicitly defined as $\widehat{p}_i(\ell) = \sum_{v\in\partial T:\ell\in S_v} 2^{-depth_T(v)}\widehat{p}_{i,v}(\ell)$.

ROUNDING: for all $v \in \partial T$ with $S_v = \{\ell_1, \ell_2\}$, $\ell_1, \ell_2 \in [k]$ do the following

1. find a set of probabilities $\{p_{i,\ell_1}\}_{i\in\mathcal{I}}$ with the following properties
   - for all $i \in \mathcal{I}$, $|p_{i,\ell_1} - p_{i,v}(\ell_1)| \le \frac{1}{z}$;
   - for all $i \in \mathcal{I}$, p_{i,ℓ_1} is an integer multiple of $\frac{1}{z}$;
   - $\left|\sum_{i\in\mathcal{I}} p_{i,\ell_1} - \sum_{i\in\mathcal{I}} p_{i,v}(\ell_1)\right| \le \frac{1}{z}$;
2. for all $i \in \mathcal{I}$, set $\widehat{p}_{i,v}(\ell_1) := p_{i,\ell_1}$, $\widehat{p}_{i,v}(\ell_2) := 1 - p_{i,\ell_1}$;

Finding the set of probabilities required by Step 1 of the ROUNDING process is straightforward and the details are omitted (see [16], Section 3.3 for a way to do so). It is now easy to check that the set of probability vectors $\{\widehat{p}_i\}_{i\in\mathcal{I}}$ satisfies Properties 1, 2 and 3 of Theorem 2.1.

3.5 Coupling within a Cell

We are now coming to the main part of the proof: Showing that the variational distance between the original and the discretized distribution within a cell depends only on z and k. We will only argue that our discretization satisfies (3); the proof of (4) is identical.

Before proceeding let us introduce some notation. Specifically,

- let $\Phi_i \in \partial T$ be the leaf chosen by Stage 1 of the process SAMPLING $\mathcal{X}_i$ and $\hat{\Phi}_i \in \partial T$ the leaf chosen by Stage 1 of SAMPLING $\widehat{\mathcal{X}}_i$;
- let $\Phi = (\Phi_i)_{i\in\mathcal{I}}$ and let G denote the distribution of Φ; similarly, let $\hat{\Phi} = (\hat{\Phi}_i)_{i\in\mathcal{I}}$ and let $\widehat{G}$ denote the distribution of $\hat{\Phi}$.

Moreover, for all $v \in \partial T$, with $S_v = \{\ell_1, \ell_2\}$ and ordering (ℓ_1, ℓ_2),

- let $\mathcal{I}_v \subseteq \mathcal{I}$ be the (random) index set such that $i \in \mathcal{I}_v$ iff $i \in \mathcal{I} \wedge \Phi_i = v$ and, similarly, let $\widehat{\mathcal{I}}_v \subseteq \mathcal{I}$ be the (random) index set such that $i \in \widehat{\mathcal{I}}_v$ iff $i \in \mathcal{I} \wedge \hat{\Phi}_i = v$;
- let $\mathcal{J}_{v,1}, \mathcal{J}_{v,2} \subseteq \mathcal{I}_v$ be the (random) index sets such $i \in \mathcal{J}_{v,1}$ iff $i \in \mathcal{I}_v \wedge \mathcal{X}_i = e_{\ell_1}$ and $i \in \mathcal{J}_{v,2}$ iff $i \in \mathcal{I}_v \wedge \mathcal{X}_i = e_{\ell_2}$;
- let $T_{v,1} = |\mathcal{J}_{v,1}|$, $T_{v,2} = |\mathcal{J}_{v,2}|$ and let F_v denote the distribution of $T_{v,1}$;
- let $T := ((T_{v,1}, T_{v,2}))_{v\in\partial T}$ and let F denote the distribution of T;
- let $\widehat{\mathcal{J}}_{v,1}, \widehat{\mathcal{J}}_{v,2}, \widehat{T}_{v,1}, \widehat{T}_{v,2}, \widehat{T}, \widehat{F}_v, \widehat{F}$ be defined similarly.

The following is easy to see; we postpone its proof to the appendix.

Claim 3.4 *For all $\theta \in (\partial T)^{\mathcal{I}}$, $G(\theta) = \widehat{G}(\theta)$.*

Since G and $\widehat{G}$ are the same distribution we will henceforth denote that distribution by G. The following lemma is sufficient to conclude the proof of Theorem 2.1.

Lemma 3.5 *There exists a value of α, used in the definition of the cells, such that, for all $v \in \partial T$,*

$$G\left(\begin{array}{l}\theta : \\ ||F_v(\cdot|\Phi=\theta) - \widehat{F}_v(\cdot|\hat{\Phi}=\theta)||_{\mathrm{TV}} \le O\left(\frac{2^k\log z}{z^{1/5}}\right)\end{array}\right) \ge 1 - \frac{4}{z^{1/3}}, \quad (5)$$

where $F_v(\cdot|\Phi)$ denotes the conditional probability distribution of $T_{v,1}$ given Φ and, similarly, $\widehat{F}_v(\cdot|\hat{\Phi})$ denotes the conditional probability distribution of $\widehat{T}_{v,1}$ given $\hat{\Phi}$.

Lemma 3.5 states roughly that, for all $v \in \partial T$, with probability at least $1 - \frac{4}{z^{1/3}}$ over the choices made by Stage 1 of processes $\{$SAMPLING $\mathcal{X}_i\}_{i\in\mathcal{I}}$ and $\{$SAMPLING $\widehat{\mathcal{X}}_i\}_{i\in\mathcal{I}}$ — assuming that these processes are coupled to make the same decisions in Stage 1 — the total variation distance between the conditional distribution of $T_{v,1}$ and $\widehat{T}_{v,1}$ is bounded by $O\left(\frac{2^k \log z}{z^{1/5}}\right)$. The following lemma, whose proof is provided in the appendix, concludes the proof of the main theorem.

Lemma 3.6 (5) *implies*

$$||F - \widehat{F}||_{\mathrm{TV}} \le O\left(k\frac{2^k \log z}{z^{1/5}}\right). \tag{6}$$

Note that (6) easily implies (3)

3.6 Proof of Lemma 3.5

To conclude the proof of Theorem 2.1, it remains to show Lemma 3.5. Roughly speaking, the proof consists of showing that, with high probability over the random walks performed in Stage 1 of SAMPLING, the one-dimensional experiment occurring at a particular leaf v of the tree is similar in both the original and the discretized distribution. The similarity is quantified by Lemmas 3.10 and 3.11 for leaves of type A and B respectively. Then, Lemmas 3.7, 3.8 and 3.9 establish that, if the experiments are sufficiently similar, they can be coupled so that their outcomes agree with high probability.

More precisely, let $v \in \partial T$, $\mathcal{S}_v = \{\ell_1, \ell_2\}$, and suppose the ordering (ℓ_1, ℓ_2). Also, let us denote $\ell_v^* = \ell_1$ and define the following functions

- $\mu_v(\theta) := \sum_{i:\theta_i = v} p_{i,v}(\ell_v^*)$;
- $\widehat{\mu}_v(\hat{\theta}) := \sum_{i:\hat{\theta}_i = v} \widehat{p}_{i,v}(\ell_v^*)$.

Note that the random variable $\mu_v(\Phi)$ represents the total probability mass that is placed on the strategy ℓ_v^* after the Stage 1 of the SAMPLING process is completed for all vectors $\mathcal{X}_i$, $i \in \mathcal{I}$. Conditioned on the outcome of Stage 1 of SAMPLING for the vectors $\{\mathcal{X}_i\}_{i\in\mathcal{I}}$, $\mu_v(\Phi)$ is the expected number of the vectors from $\mathcal{I}_v$ that will select strategy ℓ_v^* in Stage 2 of SAMPLING. Similarly, conditioned on the outcome of Stage 1 of SAMPLING for the vectors $\{\widehat{\mathcal{X}}_i\}_{i\in\mathcal{I}}$, $\widehat{\mu}_v(\widehat{\Phi})$ is the expected number of the vectors from $\widehat{\mathcal{I}}_v$ that will select strategy ℓ_v^* in Stage 2 of SAMPLING.

Intuitively, if we can couple the choices made by the random vectors $\mathcal{X}_i$, $i \in \mathcal{I}$, in Stage 1 of SAMPLING with the choices made by the random vectors $\widehat{\mathcal{X}}_i$, $i \in \mathcal{I}$, in Stage 1 of SAMPLING in such a way that, with overwhelming probability, $\mu_v(\Phi)$ and $\widehat{\mu}_v(\widehat{\Phi})$ are close, then also the conditional distributions $F_v(\cdot|\Phi)$, $\widehat{F}_v(\cdot|\hat{\Phi})$ should be close in total variation distance. The goal of this section is to make this intuition rigorous. We do this in 2 steps by showing the following.

1. The choices made in Stage 1 of SAMPLING can be coupled so that the absolute difference $|\mu_v(\Phi) - \widehat{\mu}_v(\widehat{\Phi})|$ is small with high probability. (Lemmas 3.10 and 3.11.)

2. If the absolute difference $|\mu_v(\theta) - \widehat{\mu}_v(\widehat{\theta})|$ is sufficiently small, then so is the total variation distance $||F_v(\cdot|\Phi = \theta) - \widehat{F}_v(\cdot|\hat{\Phi} = \theta)||_{TV}$. (Lemmas 3.7, 3.8, and 3.9.)

We start with Step 2 of the above program. We use different arguments depending on whether v is a Type A or Type B leaf. Let $\partial T = Ł_A \sqcup Ł_B$, where $Ł_A$ is the set of type A leaves of the cell and $Ł_B$ the set of type B leaves of the cell. For some constant β to be decided later, we show the following lemmas.

Lemma 3.7 *For some $\theta \in (\partial T)^{\mathcal{I}}$ and $v \in Ł_A$ suppose that*

$$|\mu_v(\theta) - \mathcal{E}[\mu_v(\Phi)]| \le z^{(\alpha-1)/2}\sqrt{\mathcal{E}[\mu_v(\Phi)]\log z} \tag{7}$$

$$\left|\widehat{\mu}_v(\theta) - \mathcal{E}[\widehat{\mu}_v(\hat{\Phi})]\right| \le z^{(\alpha-1)/2}\sqrt{\mathcal{E}[\widehat{\mu}_v(\hat{\Phi})]\log z} \tag{8}$$

then

$$||F_v(\cdot|\Phi = \theta) - \widehat{F}_v(\cdot|\hat{\Phi} = \theta)||_{TV} \le O\left(\frac{\sqrt{\log z}}{z^{(1-\alpha)/2}}\right).$$

Lemma 3.8 *For some $\theta \in (\partial T)^{\mathcal{I}}$ and $v \in Ł_B$ suppose that*

$$n_v(\theta) := |\{i : \theta_i = v\}| \ge z^{\beta}, \tag{9}$$

$$|\mu_v(\theta) - \widehat{\mu}_v(\theta)| \le \frac{1}{z} + \frac{\sqrt{\log z}}{z}\sqrt{|\mathcal{I}|}, \tag{10}$$

$$|n_v(\theta) - 2^{-depth_T(v)}|\mathcal{I}|| \le \sqrt{3\log z}\sqrt{2^{-depth_T(v)}|\mathcal{I}|}; \tag{11}$$

then

$$||F_v(\cdot|\Phi = \theta) - \widehat{F}_v(\cdot|\hat{\Phi} = \theta)||_{TV} \le O\left(\frac{2^{\frac{depth_T(v)}{2}}\sqrt{\log z}}{z^{\frac{1+\alpha}{2}}}\right) + O\left(\frac{2^{\frac{depth_T(v)}{2}}\log z}{z^{\frac{\alpha+\beta+1}{2}}}\right) + O(z^{-\alpha}) + O(z^{-(\frac{\alpha+\beta-1}{2})}).$$

Lemma 3.9 *For some $\theta \in (\partial T)^{\mathcal{I}}$ and $v \in Ł_B$ suppose that*

$$n_v(\theta) := |\{i : \theta_i = v\}| \le z^{\beta} \tag{12}$$

then

$$||F_v(\cdot|\Phi = \theta) - \widehat{F}_v(\cdot|\hat{\Phi} = \theta)||_{TV} \le O(z^{-(1-\beta)}).$$

The proof of Lemma 3.9 follows from a coupling argument similar to that used in the proof of Lemma 3.13 in [16] and is omitted. The proofs of Lemmas 3.7 and 3.8 can be found respectively in Sections B and C of the appendix. Lemma 3.7 provides conditions which, if satisfied by some θ at a leaf of Type A, then the conditional distributions $F_v(\cdot|\Phi = \theta)$ and $\widehat{F}_v(\cdot|\hat{\Phi} = \theta)$ are close in total variation distance. Similarly, Lemmas 3.8 and 3.9 provide conditions for the leaves of Type B. The following lemmas state that these conditions are satisfied with high probability. Their proof is given in Section D of the appendix.

Lemma 3.10 *Let $v \in Ł_A$. Then*

$$G\begin{pmatrix}\theta : |\mu_v(\theta) - \mathcal{E}[\mu_v(\Phi)]| \le \frac{\sqrt{\log z}}{z^{(1-\alpha)/2}}\sqrt{\mathcal{E}[\mu_v(\Phi)]} \\ \wedge \left|\widehat{\mu}_v(\theta) - \mathcal{E}[\widehat{\mu}_v(\hat{\Phi})]\right| \le \frac{\sqrt{\log z}}{z^{(1-\alpha)/2}}\sqrt{\mathcal{E}[\widehat{\mu}_v(\hat{\Phi})]}\end{pmatrix} \ge 1 - 4z^{-1/3}. \quad (13)$$

Lemma 3.11 *Let $v \in Ł_B$. Then*

$$G\begin{pmatrix}\theta : |\mu_v(\theta) - \widehat{\mu}_v(\theta)| \le \frac{1+\sqrt{|\mathcal{I}|\log z}}{z} \wedge \\ |n_v(\theta) - 2^{-depth_T(v)}|\mathcal{I}|| \le \sqrt{3\log z}\sqrt{2^{-depth_T(v)}|\mathcal{I}|}\end{pmatrix} \ge 1 - \frac{4}{z^{1/2}}. \quad (14)$$

Setting $\alpha = \frac{3}{5}$ and $\beta = \frac{4}{5}$, combining the above, and using that $depth_T(v) \le k$, as implied by Claim 3.3, we get (5), regardless of whether $v \in Ł_A$ or $v \in Ł_B$.

4 Extensions

Returning to our algorithm (Theorem 2.2), there are several directions in which it can be immediately generalized. To give an idea of the possibilities, let us define a *semi-anonymous game* to be a game in which

- the players are partitioned into a fixed number of *types;*
- there is another partition of the players into an arbitrary number of disjoint graphical games (see [29], games in which a node's utility depends only on its neighboring nodes) of size $O(\log n)$, where n is the total number of players, and bounded degree called *extended families*;

and the utility of each player depends on (a) his/her own strategy; (b) the overall number of other players of each type playing each strategy; and (c) it also depends, in an arbitrary way, on the strategy choices of neighboring nodes in his/her own extended family. The following result, which is only indicative of the applicability of our approach, can be shown by extending the discretization method via dynamic programming (details omitted):

Theorem 4.1 *There is a PTAS for semi-anonymous games with a fixed number of strategies.*

Further generalizations (for example, not bounding the size of the extended families) lead to PPAD-complete problems.

The discretization approach for the Nash equilibrium that we employed so far in this paper to anonymous games with a fixed number of strategies has surprisingly broad applicability, for example yielding a quasi-PTAS for general games (proof in Appendix A):

Theorem 4.2 *In any normal-form game with a constant number of strategies per player, an ϵ-approximate Nash equilibrium can be computed in time $N^{O(\log \frac{\log N}{\epsilon})}$, where N is the description size of the game.*

By combining the discretization approach with the techniques of [22] we can find approximate Nash equilibria in a large class of *graphical games*. It had long been thought that graphical games on trees with two strategies per player can be solved in polynomial time [30], until subtle flaws in the algorithm were discovered [23]. The largest class of graphical games that are known to have a polynomial-time algorithm for Nash equilibria is graphical games on a cycle and two strategies per player [23]. The following result treats a far broader class of games, albeit approximately; its proof is omitted.

Theorem 4.3 *There is PTAS for computing Nash equilibria in graphical games in which each player has a number of strategies bounded by a constant and the graph has bounded degree and $O(\log n)$ treewidth.*

5 An Application to Optimization

We illustrate an interesting application of our method in non-convex optimization. This application relates nicely to the interpretation of our main result (Theorem 2.1) as constructing a sparse cover of the set of distributions of sums of independent unit vectors under the total variation distance. The minimax optimization problem that we present arises in the context of solving repeated anonymous games, using the folk theorem [9], and similar optimization problems arise naturally in economics whenever secure strategies or threats are being computed. The optimization problem that we consider is the following.

Given functions $f_1, f_2 : \{0, 1, \ldots, n\} \to [0, 1]$ solve the optimization problem

$$\min_{p_1,\ldots,p_n \in [0,1]} \max_{k \in \{1,2\}} \left\{ \mathcal{E}_{X_i \sim B(p_i)} \left[f_k \left(\sum_{i=1}^{n} X_i \right) \right] \right\}, \tag{15}$$

where $\mathcal{E}_{X_i \sim B(p_i)}$ denotes the expectation over the joint measure of independent Bernoulli random variables $X_i, i = 1, \ldots, n$, with expectations $p_i, i = 1, \ldots, n$.

We know of no efficient algorithm for solving the above optimization problem. Nevertheless, our technique gives rise to a polynomial time approximation scheme. The idea is to use Theorem 2.1 to show that restricting the search space from $[0, 1]^n$ to $\{0, \epsilon, 2\epsilon, \ldots, 1\}^n$ results in a loss of at most $O(\epsilon^{1/6})$ in the value of the optimum. This observation is complemented by the symmetry of the objective function with respect to the p_i's; hence, we can search over the discretized space in time $O(n^{1/\epsilon})$ rather than $(1/\epsilon)^n$. The proof of the following theorem is given in detail in Appendix A.

Theorem 5.1 *There is a PTAS for solving the non-convex optimization problem* (15)*.*

The algorithm extends to the case that the minimax problem is replaced by a maximin problem. Moreover, our method provides polynomial time approximation schemes for several generalizations of (15), e.g., for the case that the maximum is taken over more than two functions (the case of one function is trivial), the domain of the functions is multidimensional (but with a constant number of dimensions), the functions have several (but constant number of) arguments, etc. Theorem 5.1 implies immediately the following result.

Corollary 5.2 ([9]) *There is a PTAS for computing threat points in repeated anonymous games with a constant number of strategies per player.*

6 Open Problems

Is there a PTAS for the Nash equilibrium problem? A major progress in this direction would be to turn the quasi-PTAS we described in the previous section for the case of a fixed number of strategies to a true PTAS. This is challenging, of course, but not hopeless. The exhaustive algorithm need not be completely exhaustive; a more intelligent search of the space, possibly in a dynamically varying grid of discretized probabilities, could possibly bring improvements in the running time. On the other hand, any constant lower bound on the approximability would be great progress as well; we conjecture that such a bound is possible at least for graphical games.

Obviously, our PTAS is not ready to be implemented and run; the exponent makes it unrealistic for any reasonable ϵ. (As we have argued in the Introduction, its true significance lies in delimiting the implications of the complexity result in [19].) There are ways to improve it, perhaps even substantially. For example, by a more elaborate trickle-down process all trees could be made full binary trees of depth $\log k$, which would remove one of the exponential functions from the exponent of the running time. But a truly practical algorithm would have to start from a new idea — possibly from that of a "less exhaustive search" mentioned in the previous paragraph. In fact, an efficient PTAS for the case of two strategies has been recently suggested [15].

References

[1] T. Abbott, D. Kane and P. Valiant. On the Complexity of Two-Player Win-Lose Games. *FOCS* 2005.

[2] A. D. Barbour. Multivariate Poisson-binomial approximation using Stein's method. *In A. D. Barbour and L. H. Y. Chen, editors, Lecture Notes Series No. 5, Institute for Mathematical Sciences, National University of Singapore,* 2005.

[3] A. D. Barbour and L. H. Y. Chen. An Introduction to Stein's Method. *In A. D. Barbour and L. H. Y. Chen (editors), Lecture Notes Series No. 4, Institute for Mathematical Sciences, National University of Singapore, Singapore University Press and World Scientific,* 2005.

[4] A. D. Barbour, L. Holst and S. Janson. Poisson Approximation. *Oxford University Press,* New York, 1992.

[5] A. D. Barbour and T. Lindvall. Translated Poisson Approximation for Markov Chains. *Journal of Theoretical Probability*, 19(3), July 2006.

[6] R. N. Bhattacharya. On Errors of Normal Approximation. *The Annals of Probability*, 3(5): 815–828, 1975.

[7] M. Blonski. Anonymous Games with Binary Actions. *Games and Economic Behavior*, 28(2): 171–180, 1999.

[8] M. Blonski. The women of Cairo: Equilibria in Large Anonymous Games. *Journal of Mathematical Economics,* 41(3): 253–264, 2005.

[9] C. Borgs, J. Chayes, N. Immorlica, A. T. Kalai, V. Mirrokni and C. H. Papadimitriou. The Myth of the Folk Theorem. *STOC* 2008.

[10] H. Bosse, J. Byrka and E. Markakis. New Algorithms for Approximate Nash Equilibria in Bimatrix Games. *WINE* 2007.

[11] X. Chen and X. Deng. Settling the Complexity of Two-Player Nash Equilibrium. *FOCS* 2006.

[12] X. Chen, X. Deng and S. Teng. Computing Nash Equilibria: Approximation and Smoothed Complexity. *FOCS* 2006.

[13] S. Chien and A. Sinclair. Convergence to Approximate Nash Equilibria in Congestion Games. *SODA* 2007.

[14] T. M. Cover and J. A. Thomas. Elements of Information Theory. *Wiley* 2006.

[15] C. Daskalakis. An Efficient PTAS for Two-Strategy Anonymous Games. *Manuscript* 2008.

[16] C. Daskalakis and C. H. Papadimitriou. Computing Equilibria in Anonymous Games. *FOCS* 2007.

[17] C. Daskalakis and C. H. Papadimitriou. Discretized Multinomial Distributions and Nash Equilibria in Anonymous Games. *ArXiv* 2008.

[18] C. Daskalakis and C. H. Papadimitriou. On Oblivious PTAS for Nash Equilibrium. *Manuscript* 2008.

[19] C. Daskalakis, P. W. Goldberg and C. H. Papadimitriou. The Complexity of Computing a Nash Equilibrium. *STOC* 2006.

[20] C. Daskalakis, A. Mehta and C. H. Papadimitriou. A Note on Approximate Nash Equilibria. *WINE* 2006.

[21] C. Daskalakis, A. Mehta and C. H. Papadimitriou. Progress in Approximate Nash Equilibria. *EC* 2007.

[22] C. Daskalakis and C. Papadimitriou. Computing Pure Nash Equilibria via Markov Random Fields. *EC* 2006.

[23] E. Elkind, L. A. Goldberg and P. W. Goldberg. Nash Equilibria in Graphical Games on Trees Revisited. *EC* 2006.

[24] T. Feder, H. Nazerzadeh and A. Saberi. Approximating Nash Equilibria Using Small-Support Strategies. *EC* 2007.

[25] D. Gale, H. W. Kuhn and A. W. Tucker. On Symmetric Games. *In H. W. Kuhn and A. W. Tucker, editors, Contributions to the Theory of Games*, 1:81–87, Princeton University Press, 1950.

[26] F. Götze. On the Rate of Convergence in the Multivariate CLT. *The Annals of Probability*, 19(2): 724–739, 1991.

[27] P. W. Goldberg and C. H. Papadimitriou. Reducibility Among Equilibrium Problems. *STOC* 2006.

[28] R. Kannan and T. Theobald. Games of Fixed Rank: A Hierarchy of Bimatrix Games. *SODA* 2007.

[29] M. Kearns, M. Littman and S. Singh. Graphical Models for Game Theory. *UAI* 2001.

[30] M. Kearns, M. Littman and S. Singh. An Efficient Exact Algorithm for Singly Connected Graphical Games. *NIPS* 2001.

[31] S. C. Kontogiannis, P. N. Panagopoulou and P. G. Spirakis. Polynomial Algorithms for Approximating Nash Equilibria of Bimatrix Games. *WINE* 2006.

[32] R. Lipton, E. Markakis and A. Mehta. Playing Large Games Using Simple Strategies. *EC* 2003.

[33] J. Nash. Noncooperative Games. *Annals of Mathematics*, 54:289–295, 1951.

[34] J. von Neumann and O. Morgenstern. Theory of Games and Economic Behavior. Princeton University Press, 1944.

[35] C. H. Papadimitriou and T. Roughgarden. Computing Equilibria in Multi-Player Games. *SODA* 2005.

[36] C. H. Papadimitriou. Computing Correlated Equilibria in Multiplayer Games. *STOC* 2005.

[37] A. Röllin. Translated Poisson Approximation Using Exchangeable Pair Couplings. *Annals of Applied Probability*, 17:1596-1614, 2007.

[38] B. Roos. Metric Multivariate Poisson Approximation of the Generalized Multinomial Distribution. *Theory of Probability and its Applications*, 43(2), 306–315, 1998.

[39] H. Tsaknakis and P. G. Spirakis. An Optimization Approach for Approximate Nash Equilibria . *WINE* 2007.

APPENDIX

A Skipped Proofs

Proof of Claim 3.3: That a tree resulting from TDP has $k-1$ leaves follows by induction: It is true when $k=2$, and for general k, the left subtree has j strategies and thus, by induction, $j-1$ leaves, and the right subtree has at most $k+1-j$ strategies and $k-j$ leaves; adding we get the result.

To estimate the number of cells, let us fix the set of strategies and their ordering at the root of the tree (thus the result of the calculation will have to be multiplied by $2^k k!$) and then count the number of trees that could be output by TDP. Suppose that the root has cardinality m and that the children of the root are assigned sets of sizes j and $m+1-j$ (or, in the event of no duplication, $m-j$), respectively. If $j=2$, then a duplication has to have happened and, for the ordering of the strategies at the left child of the root, there are at most 2 possibilities depending on whether the "divided strategy" is still the largest at the left side; similarly, for the right side there are $m-1$ possibilities: either the divided strategy is still the largest at the right side, or it is not in which case it has to be inserted at the correct place in the ordering and the last strategy of the right side must be moved to the second place. If $j>2$, similar considerations show that there are at most $j-1$ possibilities for the left side and 1 possibility for the right side. It follows that the number of trees is bounded from above by the solution $T(k)$ of the recurrence

$$\begin{aligned} T(n) &= 2\,T(2)\cdot(n-1)T(n-1) \\ &\quad + \sum_{j=3}^{n-1}(j-1)T(j)\cdot\max\{T(n-j),T(n+1-j)\}. \end{aligned}$$

with T(2)=1. It follows that the total number of trees can be upper-bounded by the function k^{k^2}. Taking into account that there are $2^k k!$ choices for the set of strategies and their ordering at the root of the tree, and that each leaf can be of either Type A, or of Type B, it follows that the total number of cells is bounded by $g(k)=k^{k^2}2^{k-1}2^k k!$. ■

Proof of Claim 3.4: The proof follows by a straightforward coupling argument. Indeed, for all $i\in\mathcal{I}$, let us couple the choices made by Stage 1 of SAMPLING $\mathcal{X}_i$ and SAMPLING $\widehat{\mathcal{X}}_i$ so that the random leaf $\Phi_i\in\partial T$ chosen by SAMPLING $\mathcal{X}_i$ and the random leaf $\hat{\Phi}_i\in\partial T$ chosen by SAMPLING $\widehat{\mathcal{X}}_i$ are equal, that is, for all $i\in\mathcal{I}$, in the joint probability space $\Pr[\Phi_i=\hat{\Phi}_i]=1$; the existence of such a coupling is straightforward since Stage 1 of both SAMPLING $\mathcal{X}_i$ and SAMPLING $\widehat{\mathcal{X}}_i$ is the same random walk on T. ■

Proof of Lemma 3.6: Note first that (5) implies via a union bound that

$$G\left(\theta:\forall v\in\partial T, ||F_v(\cdot|\Phi=\theta)-\widehat{F}_v(\cdot|\hat{\Phi}=\theta)||_{TV}\le O\left(\frac{2^k\log z}{z^{1/5}}\right)\right)\ge 1-O(kz^{-1/3}), \tag{16}$$

since, by Claim 3.3, the number of leaves is at most $k-1$.

Now suppose that, for a given value of $\theta\in(\partial T)^{\mathcal{I}}$, the following is satisfied

$$\forall v\in\partial T,\ ||F_v(\cdot|\Phi=\theta)-\widehat{F}_v(\cdot|\hat{\Phi}=\theta)||_{TV}\le O\left(\frac{2^k\log z}{z^{1/5}}\right). \tag{17}$$

Observe that the variables $\{T_{v,1}\}_{v\in\partial T}$ are conditionally independent given Φ, and, similarly, the variables $\{\widehat{T}_{v,1}\}_{v\in\partial T}$ are conditionally independent given $\hat{\Phi}$. This, by the coupling lemma, Claim 3.3, and (17) implies that

$$||F(\cdot|\Phi=\theta)-\widehat{F}(\cdot|\hat{\Phi}=\theta)||_{TV}\le O\left(k\frac{2^k\log z}{z^{1/5}}\right),$$

where we used that, if $\Phi=\hat{\Phi}$, then $|\mathcal{I}_v|=|\widehat{\mathcal{I}}_v|$, $\forall v\in\partial T$.

Therefore, (16) implies

$$G\left(\theta:||F(\cdot|\Phi=\theta)-\widehat{F}(\cdot|\hat{\Phi}=\theta)||_{TV}\le O\left(k\frac{2^k\log z}{z^{1/5}}\right)\right)\ge 1-O(kz^{-1/3}). \tag{18}$$

All that remains is to shift the bound of (18) to the unconditional space. The following lemma establishes this reduction.

Lemma A.1 (18) *implies*

$$||F - \widehat{F}||_{\mathrm{TV}} \leq O\left(k \frac{2^k \log z}{z^{1/5}}\right). \tag{19}$$

Proof of Lemma A.1: Let us denote by

$$Good = \{\theta | \theta \in (\partial T)^{\mathcal{I}} : ||F(\cdot|\Phi = \theta) - \widehat{F}(\cdot|\hat{\Phi} = \theta)||_{TV} \leq O\left(k \frac{2^k \log z}{z^{1/5}}\right),$$

and let $Bad = (\partial T)^{\mathcal{I}} - Good$. By (18), if follows that $G(Bad) \leq O(kz^{-1/3})$.

$$\begin{aligned}
||T - \widehat{T}||_{TV} &= \frac{1}{2}\sum_t |F(t) - \widehat{F}(t)| \\
&= \frac{1}{2}\sum_t \left|\sum_\theta F(t|\Phi = \theta)G(\Phi = \theta) - \sum_\theta \widehat{F}(t|\hat{\Phi} = \theta)\widehat{G}(\hat{\Phi} = \theta)\right| \\
&= \frac{1}{2}\sum_t \left|\sum_\theta (F(t|\Phi = \theta) - \widehat{F}(t|\hat{\Phi} = \theta))G(\theta)\right| \qquad \left(\text{using } G(\theta) = \widehat{G}(\theta), \forall\theta\right) \\
&\leq \frac{1}{2}\sum_t \sum_\theta \left|F(t|\Phi = \theta) - \widehat{F}(t|\hat{\Phi} = \theta)\right| G(\theta) \\
&= \frac{1}{2}\sum_t \sum_{\theta \in Good} \left|F(t|\Phi = \theta) - \widehat{F}(t|\hat{\Phi} = \theta)\right| G(\theta) \\
&\quad + \frac{1}{2}\sum_t \sum_{\theta \in Bad} \left|F(t|\Phi = \theta) - \widehat{F}(t|\hat{\Phi} = \theta)\right| G(\theta) \\
&\leq \sum_{\theta \in Good} G(\theta)\left(\frac{1}{2}\sum_t \left|F(t|\Phi = \theta) - \widehat{F}(t|\hat{\Phi} = \theta)\right|\right) \\
&\quad + \sum_{\theta \in Bad} G(\theta)\left(\frac{1}{2}\sum_t \left|F(t|\Phi = \theta) - \widehat{F}(t|\hat{\Phi} = \theta)\right|\right) \\
&\leq \sum_{\theta \in Good} G(\theta) \cdot O\left(k \frac{2^k \log z}{z^{1/5}}\right) + \sum_{\theta \in Bad} G(\theta) \\
&\leq O\left(k \frac{2^k \log z}{z^{1/5}}\right) + O(kz^{-1/3}).
\end{aligned}$$

■

■

Proof of Theorem 4.2: Let p be the number of players and s the number of strategies per player which we assume to be a constant; the input size is $N = ps^p$. Consider a new p-player game in which the set of pure strategies of each player is the set of all distributions over the s strategies of the original game whose probabilities are integer multiples of $\delta = \frac{\epsilon}{2ps}$. We claim that, if we search over all the pure strategy profiles of the new game, we are bound to discover an ϵ-approximate Nash equilibrium of the original game. To prove this, it suffices to notice the following which is proven below by two applications of the coupling lemma.

Lemma A.2 *Let $(x_1, \ldots, x_p)$ be a mixed-Nash equilibrium of the original game and $(\hat{x}_1, \ldots, \hat{x}_p)$ be another set of mixed strategies, where, for all i, j, $\hat{x}_i(j)$ is an integer multiple of $\delta = \frac{\epsilon}{2ps}$, $|\hat{x}_i(j) - x_i(j)| \leq \delta$ and, if $x_i(j) = 0$, then also $\hat{x}_i(j)$=0. Then $(\hat{x}_1, \ldots, \hat{x}_p)$ is an ϵ-approximate Nash equilibrium of the original game.*

The number of pure strategy profiles of the new game that we have to search over is at most $\left(\left(\frac{1}{\delta}\right)^s\right)^p$, which is easily seen to be $N^{O(\log \frac{\log N}{\epsilon})}$. ■

Proof of Lemma A.2: For every player i and strategy j, let U_j^i and $\hat{U}_j^i$ be the expected utility of player i if she plays j and the other players play $\{x_{i'}\}_{i' \neq i}$ and $\{\hat{x}_{i'}\}_{i' \neq i}$ respectively. The difference between U_j^i and $\hat{U}_j^i$ can be bounded as follows

$$|U_j^i - \hat{U}_j^i| \leq ||(x_1, \ldots, x_p) - (\hat{x}_1, \ldots, \hat{x}_p)||_{TV},$$

where the right hand side of the above expression represents the total variation distance between the compound distributions $(x_1, \ldots, x_p)$ and $(\hat{x}_1, \ldots, \hat{x}_p)$, and we used the fact that the payoff functions of the players lie in $[0, 1]$. We will show that

$$||(x_1, \ldots, x_p) - (\hat{x}_1, \ldots, \hat{x}_p)||_{TV} \leq \frac{\epsilon}{2}.$$

Indeed, for all i, let $\mathcal{X}_i$ be a random s-dimensional vector such that $\mathcal{X}_i = e_j$ with probability $x_i(j)$, and suppose that the vectors $\{\mathcal{X}_i\}_i$ are independent. Similarly, define vectors $\{\hat{\mathcal{X}}_i\}_i$. The coupling lemma implies that, for any coupling of $\{\mathcal{X}_i\}_i$ and $\{\hat{\mathcal{X}}_i\}_i$,

$$||(\mathcal{X}_1, \ldots, \mathcal{X}_p) - (\hat{\mathcal{X}}_1, \ldots, \hat{\mathcal{X}}_p)||_{TV} \leq \Pr[(\mathcal{X}_1, \ldots, \mathcal{X}_p) \neq (\hat{\mathcal{X}}_1, \ldots, \hat{\mathcal{X}}_p)],$$

which, by a union bound, implies

$$||(\mathcal{X}_1, \ldots, \mathcal{X}_p) - (\hat{\mathcal{X}}_1, \ldots, \hat{\mathcal{X}}_p)||_{TV} \leq \sum_i \Pr[\mathcal{X}_i \neq \hat{\mathcal{X}}_i].$$

Let us now fix a coupling for which, for all i,

$$\Pr[\mathcal{X}_i \neq \hat{\mathcal{X}}_i] = ||\mathcal{X}_i - \hat{\mathcal{X}}_i||_{TV}.$$

Such a coupling exists by the coupling lemma and the fact that the random vectors $\{\mathcal{X}_i\}_i$ are independent and so are the random vectors $\{\hat{\mathcal{X}}_i\}_i$. Combining the above, we get

$$||(\mathcal{X}_1, \ldots, \mathcal{X}_p) - (\hat{\mathcal{X}}_1, \ldots, \hat{\mathcal{X}}_p)||_{TV} \leq \sum_i ||\mathcal{X}_i - \hat{\mathcal{X}}_i||_{TV}.$$

Observe finally that

$$||(x_1, \ldots, x_p) - (\hat{x}_1, \ldots, \hat{x}_p)||_{TV} = ||(\mathcal{X}_1, \ldots, \mathcal{X}_p) - (\hat{\mathcal{X}}_1, \ldots, \hat{\mathcal{X}}_p)||_{TV},$$

and, for all i,

$$||\mathcal{X}_i - \hat{\mathcal{X}}_i||_{TV} = ||x_i - \hat{x}_i||_{TV} \leq \delta s = \frac{\epsilon}{2p}.$$

It follows that

$$||(x_1, \ldots, x_p) - (\hat{x}_1, \ldots, \hat{x}_p)||_{TV} \leq \frac{\epsilon}{2}.$$

Hence, for all i, j,

$$|U_j^i - \hat{U}_j^i| \leq ||(x_1, \ldots, x_p) - (\hat{x}_1, \ldots, \hat{x}_p)||_{TV} \leq \frac{\epsilon}{2},$$

which implies that $(\hat{x}_1, \ldots, \hat{x}_p)$ is an ϵ-approximate Nash equilibrium of the original game. ■

Proof of Theorem 5.1: It is not hard to see that for any sets of probabilities $\{p_i\}_i$ and $\{p'_i\}_i$, and for any $\alpha \in \{1, 2\}$,

$$\left| \mathcal{E}_{X_i \sim B(p_i)} \left[f_\alpha \left(\sum_{i=1}^n X_i \right) \right] - \mathcal{E}_{Y_i \sim B(p'_i)} \left[f_\alpha \left(\sum_{i=1}^n Y_i \right) \right] \right| \leq \left|\left| \sum_i X_i - \sum_i Y_i \right|\right|_{TV},$$

where, in the right hand side of the above $\{X_i\}_i$ is a set of independent Bernoulli random variables with expectations $\{p_i\}_i$ and $\{Y_i\}_i$ a set of independent Bernoulli random variables with expectations $\{p'_i\}_i$.

Suppose now that $\{p_i^*\}_i$ is the set of probabilities achieving the optimum value for (15). It follows from the above that if we perturb the p_i^*'s to another set of probabilities $\{p_i'^*\}_i$ the value of the minmax problem is only affected by an additive term $||\sum_i X_i - \sum_i Y_i||_{TV}$, where $X_i \sim B(p_i^*)$ and $Y_i \sim B(p_i'^*)$, for all i.

It follows from Theorem 2.1 that, for any set of probabilities $\{p_i^*\}_i$, there exists another set of ϵ-"discretized" probabilities $\{p_i'^*\}_i$, that is, $p_i'^*$ is an integer multiple of ϵ, for all i, such that

$$\left\|\sum_i X_i - \sum_i Y_i\right\|_{TV} \leq O(\epsilon^{1/6}).$$

Hence, we can restrict the optimization to ϵ-discretized probabilities with an additive loss of $O(\epsilon^{1/6})$ in the value of the optimum. Even so, the search space is of size $\Omega\left(\left(\frac{1}{\epsilon}\right)^n\right)$ which is exponential in the input size $O(n)$. By observing that the objective function is symmetric with respect to the set of probabilities $\{p_i\}_i$ we can prune the search space to searching only over the partitions of n unlabeled objects into $1/\epsilon$ bins, that is $O(n^{1/\epsilon})$ possible partitions. This results in a polynomial time approximation scheme. ■

B Proof of Lemma 3.7

Proof: By the assumption it follows that

$$|\mu_v(\theta) - \widehat{\mu}_v(\theta)| \leq \left|\mathcal{E}[\mu_v(\Phi)] - \mathcal{E}[\widehat{\mu}_v(\hat{\Phi})]\right| + z^{(\alpha-1)/2}\sqrt{\mathcal{E}[\mu_v(\Phi)]\log z} + z^{(\alpha-1)/2}\sqrt{\mathcal{E}[\widehat{\mu}_v(\hat{\Phi})]\log z}.$$

Moreover, note that

$$\mathcal{E}[\mu_v(\Phi)] = 2^{-depth_T(v)}\sum_{i\in\mathcal{I}} p_{i,v}(\ell_v^*)$$

and, similarly,

$$\mathcal{E}[\widehat{\mu}_v(\hat{\Phi})] = 2^{-depth_T(v)}\sum_{i\in\mathcal{I}} \widehat{p}_{i,v}(\ell_v^*).$$

By the definition of the ROUNDING procedure it follows that

$$|\mathcal{E}[\mu_v(\Phi)] - \mathcal{E}[\widehat{\mu}_v(\hat{\Phi})]| \leq 2^{-depth_T(v)}\frac{1}{z}.$$

Hence it follows that

$$|\mu_v(\theta) - \widehat{\mu}_v(\theta)| \leq 2^{-depth_T(v)}\frac{1}{z} + \frac{2\sqrt{\log z}}{z^{(1-\alpha)/2}}\sqrt{\max\{\mathcal{E}[\mu_v(\Phi), \mathcal{E}[\widehat{\mu}_v(\hat{\Phi})]\}]}. \quad (20)$$

Let $\mathcal{N}_v(\theta) := \{i : \theta_i = v\}$, $n_v = |\mathcal{N}_v|$. Conditioned on $\Phi = \theta$, the distribution of $T_{v,1}$ is the sum of n_v independent Bernoulli random variables $\{Z_i\}_{i\in\mathcal{N}_v}$ with expectations $\mathcal{E}[Z_i] = p_{i,v}(\ell_v^*) \leq \frac{\lfloor z^\alpha \rfloor}{z}$. Similarly, conditioned on $\hat{\Phi} = \theta$, the distribution of $\widehat{T}_{v,1}$ is the sum of n_v independent Bernoulli random variables $\{\widehat{Z}_i\}_{i\in\mathcal{N}_v}$ with expectations $\mathcal{E}[\widehat{Z}_i] = \widehat{p}_{i,v}(\ell_v^*) \leq \frac{\lfloor z^\alpha \rfloor}{z}$. Note that

$$\mathcal{E}\left[\sum_{i\in\mathcal{N}_v} Z_i\right] = \mu_v(\theta)$$

and, similarly,

$$\mathcal{E}\left[\sum_{i\in\mathcal{N}_v} \widehat{Z}_i\right] = \widehat{\mu}_v(\theta).$$

Without loss of generality, let us assume that $\mathcal{E}[\mu_v(\Phi)] \geq \mathcal{E}[\widehat{\mu}_v(\hat{\Phi})]$. Let us further distinguish two cases for some constant $\tau < 1-\alpha$ to be decided later

Case 1: $\mathcal{E}[\mu_v(\Phi)] \leq \frac{1}{z^\tau}$.

From (7) it follows that,

$$\mu_v(\theta) \le \mathcal{E}[\mu_v(\Phi)] + z^{(\alpha-1)/2}\sqrt{\mathcal{E}[\mu_v(\Phi)]\log z} \le \frac{1}{z^\tau} + \frac{\sqrt{\log z}}{z^{(\tau+1-\alpha)/2}} =: g(z).$$

Similarly, because $\mathcal{E}[\widehat{\mu}_v(\hat{\Phi})] \le \mathcal{E}[\mu_v(\Phi)] \le \frac{1}{z^\tau}$, $\widehat{\mu}_v(\theta) \le g(z)$.

By Markov's inequality, $\Pr_{\Phi=\theta}[\sum_{i\in\mathcal{N}_v} Z_i \ge 1] \le \frac{\mu_v(\theta)}{1} \le g(z)$ and, similarly, $\Pr_{\hat{\Phi}=\theta}[\sum_{i\in\mathcal{N}_v} \widehat{Z}_i \ge 1] \le g(z)$. Hence,

$$\begin{aligned}\left|\Pr{}_{\Phi=\theta}\left[\sum_{i\in\mathcal{N}_v} Z_i = 0\right] - \Pr{}_{\hat{\Phi}=\theta}\left[\sum_{i\in\mathcal{N}_v} \widehat{Z}_i = 0\right]\right| &= \left|\Pr{}_{\Phi=\theta}\left[\sum_{i\in\mathcal{N}_v} Z_i \ge 1\right] - \Pr{}_{\hat{\Phi}=\theta}\left[\sum_{i\in\mathcal{N}_v} \widehat{Z}_i \ge 1\right]\right| \\ &\le 2g(z).\end{aligned}$$

It follows then easily that

$$||F_v(\cdot|\Phi=\theta) - \widehat{F}_v(\cdot|\hat{\Phi}=\theta)||_{TV} \le 4g(z) = 4\cdot\left(\frac{1}{z^\tau} + \frac{\sqrt{\log z}}{z^{(\tau+1-\alpha)/2}}\right). \tag{21}$$

Case 2: $\mathcal{E}[\mu_v(\Phi)] \ge \frac{1}{z^\tau}$.

The following claim was proven in [16], Lemma 3.9,

Claim B.1 *For any set of independent Bernoulli random variables $\{Z_i\}_i$ with expectations $\mathcal{E}[Z_i] \le \frac{\lfloor z^\alpha\rfloor}{z}$,*

$$\left\|\sum_i Z_i - Poisson\left(\mathcal{E}\left(\sum_i Z_i\right)\right)\right\|_{\mathrm{TV}} \le \frac{1}{z^{1-\alpha}}.$$

By application of this lemma it follows that

$$\left\|\sum_{i\in\mathcal{N}_v} Z_i - Poisson(\mu_v(\theta))\right\|_{TV} \le \frac{1}{z^{1-\alpha}}, \tag{22}$$

$$\left\|\sum_{i\in\mathcal{N}_v} \widehat{Z}_i - Poisson(\widehat{\mu}_v(\theta))\right\|_{TV} \le \frac{1}{z^{1-\alpha}}. \tag{23}$$

We study next the distance between the two Poisson distributions. We use the following lemma whose proof is postponed till later in this section.

Lemma B.2 *If $\lambda = \lambda_0 + D$ for some $D > 0$, $\lambda_0 > 0$,*

$$\|Poisson(\lambda) - Poisson(\lambda_0)\|_{\mathrm{TV}} \le D\sqrt{\frac{2}{\lambda_0}}.$$

An application of Lemma B.2 gives

$$\|Poisson(\mu_v(\theta)) - Poisson(\widehat{\mu}_v(\theta))\|_{TV} \le |\mu_v(\theta) - \widehat{\mu}_v(\theta)|\sqrt{\frac{2}{\min\{\mu_v(\theta), \widehat{\mu}_v(\theta)\}}}. \tag{24}$$

We conclude with the following lemma proved in the end of this section.

Lemma B.3 *From* (7), (8), (20) *and the assumption $\mathcal{E}[\mu_v(\Phi)] \ge \frac{1}{z^\tau}$, it follows that*

$$|\mu_v(\theta) - \widehat{\mu}_v(\theta)|\sqrt{\frac{2}{\min\{\mu_v(\theta), \widehat{\mu}_v(\theta)\}}} \le \sqrt{72\frac{\log z}{z^{1-\alpha}}}.$$

Combining (22), (23), (24) and Lemma B.3 we get

$$\left\| \sum_{i \in \mathcal{N}_v} Z_i - \sum_{i \in \mathcal{N}_v} \widehat{Z}_i \right\|_{TV} \leq \frac{2}{z^{1-\alpha}} + \sqrt{72 \frac{\log z}{z^{1-\alpha}}} = O\left(\frac{\sqrt{\log z}}{z^{(1-\alpha)/2}} \right),$$

which implies

$$||F_v(\cdot|\Phi = \theta) - \widehat{F}_v(\cdot|\hat{\Phi} = \theta)||_{TV} \leq O\left(\frac{\sqrt{\log z}}{z^{(1-\alpha)/2}} \right). \tag{25}$$

Taking $\tau > (1-\alpha)/2$, we get from (21), (25) that in both cases

$$||F_v(\cdot|\Phi = \theta) - \widehat{F}_v(\cdot|\hat{\Phi} = \theta)||_{TV} \leq O\left(\frac{\sqrt{\log z}}{z^{(1-\alpha)/2}} \right). \tag{26}$$

■

Proof of lemma B.2: We make use of the following lemmas.

Lemma B.4 *If $\lambda, \lambda_0 > 0$, the Kullback-Leibler divergence between $Poisson(\lambda_0)$ and $Poisson(\lambda)$ is given by*

$$\Delta_{KL}(Poisson(\lambda)||Poisson(\lambda_0)) = \lambda \left(1 - \frac{\lambda_0}{\lambda} + \frac{\lambda_0}{\lambda} \log \frac{\lambda_0}{\lambda} \right).$$

Lemma B.5 (e.g. [14]) *If P and Q are probability measures on the same measure space and P is absolutely continuous with respect to Q then*

$$\|P - Q\|_{\mathrm{TV}} \leq \sqrt{2\Delta_{KL}(P||Q)}.$$

By simple calculus we have that

$$\Delta_{KL}(Poisson(\lambda)||Poisson(\lambda_0)) = \lambda \left(1 - \frac{\lambda_0}{\lambda} + \frac{\lambda_0}{\lambda} \log \frac{\lambda_0}{\lambda} \right) \leq \frac{D^2}{\lambda_0}.$$

Then by Lemma B.5 it follows that

$$\|Poisson(\lambda) - Poisson(\lambda_0)\|_{TV} \leq D\sqrt{\frac{2}{\lambda_0}}.$$

■

Proof of lemma B.3: From (20) and the assumption $\mathcal{E}[\mu_v(\Phi)] \geq \mathcal{E}[\widehat{\mu}_v(\hat{\Phi})]$ we have

$$|\mu_v(\theta) - \widehat{\mu}_v(\theta)|^2 \leq \frac{1}{z^2} + \frac{4\log z}{z^{1-\alpha}} \mathcal{E}[\mu_v(\Phi)] + 4\frac{1}{z}\frac{\sqrt{\log z}}{z^{(1-\alpha)/2}} \sqrt{\mathcal{E}[\mu_v(\Phi)]}.$$

From the assumption $\mathcal{E}[\mu_v(\Phi)] \geq \frac{1}{z^\tau}$ it follows

$$\begin{aligned} \mathcal{E}[\mu_v(\Phi)] &= \sqrt{\mathcal{E}[\mu_v(\Phi)]}\sqrt{\mathcal{E}[\mu_v(\Phi)]} && (27) \\ &\geq \frac{1}{z^{\tau/2}} \sqrt{\mathcal{E}[\mu_v(\Phi)]}. && (28) \end{aligned}$$

Since $\tau < 1 - \alpha$, it follows that, for sufficiently large z which only depends on α and τ, $\frac{1}{z^{\tau/2}} \geq \frac{2\sqrt{\log z}}{z^{(1-\alpha)/2}}$. Hence,

$$\mathcal{E}[\mu_v(\Phi)] \geq \frac{2\sqrt{\log z}}{z^{(1-\alpha)/2}} \sqrt{\mathcal{E}[\mu_v(\Phi)]},$$

which together with (7) implies

$$\mu_v(\theta) \geq \mathcal{E}[\mu_v(\Phi)] - z^{(\alpha-1)/2}\sqrt{\mathcal{E}[\mu_v(\Phi)]\log z} \geq \frac{1}{2}\mathcal{E}[\mu_v(\Phi)] \tag{29}$$

Similarly, starting from $\mathcal{E}[\widehat{\mu}_v(\hat{\Phi})] \geq \mathcal{E}[\mu_v(\Phi)] - \frac{1}{z} \geq \frac{1}{z^\tau} - \frac{1}{z}$, it can be shown that for sufficiently large z

$$\widehat{\mu}_v(\theta) \geq \frac{1}{2}\mathcal{E}[\widehat{\mu}_v(\hat{\Phi})]. \tag{30}$$

From (29), (30) it follows that

$$\min\{\mu_v(\theta), \widehat{\mu}_v(\theta)\} \geq \frac{1}{2}\min\{\mathcal{E}[\mu_v(\Phi)], \mathcal{E}[\widehat{\mu}_v(\hat{\Phi})]\} = \frac{1}{2}\mathcal{E}[\widehat{\mu}_v(\Phi)] \geq \frac{1}{2}\mathcal{E}[\mu_v(\Phi)] - \frac{1}{2z} \geq \frac{1}{4}\mathcal{E}[\mu_v(\Phi)],$$

where we used that $\mathcal{E}[\mu_v(\Phi)] \geq \frac{1}{z^\tau} \geq \frac{2}{z}$ for sufficiently large z, since $\tau < 1-\alpha$. Combining the above we get

$$\begin{aligned}
\frac{2|\mu_v(\theta) - \widehat{\mu}_v(\theta)|^2}{\min\{\mu_v(\theta), \widehat{\mu}_v(\theta)\}} &\leq 2\frac{\frac{1}{z^2} + \frac{4\log z}{z^{1-\alpha}}\mathcal{E}[\mu_v(\Phi)] + 4\frac{1}{z}\frac{\sqrt{\log z}}{z^{(1-\alpha)/2}}\sqrt{\mathcal{E}[\mu_v(\Phi)]}}{\frac{1}{4}\mathcal{E}[\mu_v(\Phi)]} \\
&\leq 8\frac{1}{z^2\mathcal{E}[\mu_v(\Phi)]} + 32\frac{\log z}{z^{1-\alpha}} + 32\frac{\sqrt{\log z}}{z^{1+(1-\alpha)/2}\sqrt{\mathcal{E}[\mu_v(\Phi)]}} \\
&\leq 8\frac{z^\tau}{z^2} + 32\frac{\log z}{z^{1-\alpha}} + 32\frac{z^{\tau/2}\sqrt{\log z}}{z^{1+(1-\alpha)/2}} \\
&\leq 8\frac{1}{z^{2-\tau}} + 32\frac{\log z}{z^{1-\alpha}} + 32\frac{\sqrt{\log z}}{z^{(3-\alpha-\tau)/2}} \\
&\leq 72\frac{\log z}{z^{1-\alpha}},
\end{aligned}$$

since $2-\tau > 1-\alpha$ and $(3-\alpha-\tau)/2 > 1-\alpha$, assuming sufficiently large z. ■

C Proof of Lemma 3.8

Proof: We will derive our bound by approximating with the *translated Poisson distribution*, which is defined next.

Definition C.1 ([37]) We say that an integer random variable Y has a *translated Poisson distribution* with paremeters μ and σ^2 and write

$$Ł(Y) = TP(\mu, \sigma^2)$$

if $Ł(Y - \lfloor \mu - \sigma^2 \rfloor) = Poisson(\sigma^2 + \{\mu - \sigma^2\})$, where $\{\mu - \sigma^2\}$ represents the fractional part of $\mu - \sigma^2$.

The following lemma provides a bound for the total variation distance between two translated Poisson distributions with different parameters.

Lemma C.2 ([5]) *Let $\mu_1, \mu_2 \in \mathbb{R}$ and $\sigma_1^2, \sigma_2^2 \in \mathbb{R}_+ \setminus \{0\}$ be such that $\lfloor \mu_1 - \sigma_1^2 \rfloor \leq \lfloor \mu_2 - \sigma_2^2 \rfloor$. Then*

$$\left|\left|TP(\mu_1, \sigma_1^2) - TP(\mu_2, \sigma_2^2)\right|\right|_{\mathrm{TV}} \leq \frac{|\mu_1 - \mu_2|}{\sigma_1} + \frac{|\sigma_1^2 - \sigma_2^2| + 1}{\sigma_1^2}.$$

The following lemma was proven in [16], Lemma 3.14,

Lemma C.3 *Let $z > 0$ be some integer and $\{Z_i\}_{i=1}^m$, where $m \geq z^\beta$, be any set of independent Bernoulli random variables with expectations $\mathcal{E}[Z_i] \in \left[\frac{\lfloor z^\alpha \rfloor}{z}, \frac{1}{2}\right]$. Let $\mu_1 = \sum_{i=1}^m \mathcal{E}[Z_i]$ and $\sigma_1^2 = \sum_{i=1}^m \mathcal{E}[Z_i](1 - \mathcal{E}[Z_i])$. Then*

$$\left\|\sum_{i=1}^m Z_i - TP\left(\mu_1, \sigma_1^2\right)\right\|_{\mathrm{TV}} \leq O\left(z^{-\frac{\alpha+\beta-1}{2}}\right).$$

Let $\mathcal{N}_v(\theta) := \{i : \theta_i = v\}$, $n_v(\theta) = |\mathcal{N}_v(\theta)|$. Conditioned on $\Phi = \theta$, the distribution of $T_{v,1}$ is the sum of $n_v(\theta)$ independent Bernoulli random variables $\{Z_i\}_{i\in\mathcal{N}_v(\theta)}$ with expectations $\mathcal{E}[Z_i] = p_{i,v}(\ell_v^*)$. Similarly, conditioned on $\hat{\Phi} = \theta$, the distribution of $\widehat{T}_{v,1}$ is the sum of $n_v(\theta)$ independent Bernoulli random variables $\{\widehat{Z}_i\}_{i\in\mathcal{N}_v(\theta)}$ with expectations $\mathcal{E}[\widehat{Z}_i] = \widehat{p}_{i,v}(\ell_v^*)$. Note that

$$\sum_{i\in\mathcal{N}_v(\theta)} \mathcal{E}\left[Z_i\right] = \mu_v(\theta)$$

and, similarly,

$$\sum_{i\in\mathcal{N}_v(\theta)} \mathcal{E}\left[\widehat{Z}_i\right] = \widehat{\mu}_v(\theta).$$

Setting $\mu_1 := \mu_v(\theta)$, $\mu_2 := \widehat{\mu}_v(\theta)$ and

$$\sigma_1^2 = \sum_{i\in\mathcal{N}_v(\theta)} \mathcal{E}\left[Z_i\right](1 - \mathcal{E}\left[Z_i\right]),$$

$$\sigma_2^2 = \sum_{i\in\mathcal{N}_v(\theta)} \mathcal{E}\left[\widehat{Z}_i\right](1 - \mathcal{E}\left[\widehat{Z}_i\right]),$$

we have from Lemma C.3 that

$$\left\|\sum_{i\in\mathcal{N}_v(\theta)} Z_i - TP\left(\mu_1, \sigma_1^2\right)\right\|_{TV} \leq O\left(z^{-\frac{\alpha+\beta-1}{2}}\right). \tag{31}$$

$$\left\|\sum_{i\in\mathcal{N}_v(\theta)} \widehat{Z}_i - TP\left(\mu_2, \sigma_2^2\right)\right\|_{TV} \leq O\left(z^{-\frac{\alpha+\beta-1}{2}}\right). \tag{32}$$

It remains to bound the total variation distance between the translated poisson distributions using Lemma C.2. Without loss of generality let us assume $\lfloor \mu_1 - \sigma_1^2 \rfloor \leq \lfloor \mu_2 - \sigma_2^2 \rfloor$. Note that

$$\sigma_1^2 = \sum_{i\in\mathcal{N}_v(\theta)} \mathcal{E}\left[Z_i\right](1 - \mathcal{E}\left[Z_i\right]) \geq n_v(\theta)\frac{\lfloor z^\alpha \rfloor}{z}\left(1 - \frac{\lfloor z^\alpha \rfloor}{z}\right) \geq \frac{1}{2}n_v(\theta)\frac{\lfloor z^\alpha \rfloor}{z},$$

where the last inequality holds for values of z which are larger than some function of constant α. Also,

$$\begin{aligned}
|\sigma_1^2 - \sigma_2^2| &\leq \sum_{i\in\mathcal{N}_v(\theta)} \left|\mathcal{E}\left[Z_i\right](1 - \mathcal{E}\left[Z_i\right]) - \mathcal{E}\left[\widehat{Z}_i\right](1 - \mathcal{E}\left[\widehat{Z}_i\right])\right| \\
&= \sum_{i\in\mathcal{N}_v(\theta)} \left|p_{i,v}(\ell_v^*)(1 - p_{i,v}(\ell_v^*)) - \widehat{p}_{i,v}(\ell_v^*)(1 - \widehat{p}_{i,v}(\ell_v^*))\right| \\
&= \sum_{i\in\mathcal{N}_v(\theta)} \left(|p_{i,v}(\ell_v^*) - \widehat{p}_{i,v}(\ell_v^*)| + \left|p_{i,v}^2(\ell_v^*) - \widehat{p}_{i,v}^2(\ell_v^*)\right|\right) \\
&\leq \sum_{i\in\mathcal{N}_v(\theta)} \frac{3}{z} \qquad\qquad \left(\text{using } |p_{i,v}(\ell_v^*) - \widehat{p}_{i,v}(\ell_v^*)| \leq \frac{1}{z}\right) \\
&\leq \frac{3n_v(\theta)}{z}.
\end{aligned}$$

Using the above and Lemma C.2 we have that

$$\begin{aligned} \big|\big|TP(\mu_1,\sigma_1^2)-TP(\mu_2,\sigma_2^2)\big|\big| &\le \frac{|\mu_1-\mu_2|}{\sigma_1}+\frac{|\sigma_1^2-\sigma_2^2|}{\sigma_1^2}+\frac{1}{\sigma_1^2} \\ &\le \frac{|\mu_1-\mu_2|}{\sigma_1}+\frac{\frac{3n_v(\theta)}{z}}{\frac{1}{2}n_v(\theta)\frac{\lfloor z^\alpha\rfloor}{z}}+\frac{1}{\frac{1}{2}n_v(\theta)\frac{\lfloor z^\alpha\rfloor}{z}} \\ &\le \frac{|\mu_1-\mu_2|}{\sigma_1}+O(z^{-\alpha})+\frac{1}{\frac{1}{2}z^\beta\frac{\lfloor z^\alpha\rfloor}{z}} \\ &\le \frac{|\mu_1-\mu_2|}{\sigma_1}+O(z^{-\alpha})+O(z^{-(\alpha+\beta-1)}). \end{aligned}$$

To bound the ratio $\frac{|\mu_1-\mu_2|}{\sigma_1}$ we distinguish the following cases:

- $\sqrt{3\log z}\sqrt{2^{-depth_T(v)}}\sqrt{|\mathcal{I}|}\le \frac{1}{2}2^{-depth_T(v)}|\mathcal{I}|$: Combining this inequality with (11) we get that

$$|\mathcal{I}|\le 2^{1+depth_T(v)}n_v(\theta).$$

  Hence,

$$\frac{|\mu_1-\mu_2|}{\sigma_1}\le\frac{\frac{1}{z}+\frac{\sqrt{\log z}}{z}\sqrt{|\mathcal{I}|}}{\sqrt{\frac{1}{2}n_v(\theta)\frac{\lfloor z^\alpha\rfloor}{z}}}\le\frac{\frac{1}{z}+\frac{\sqrt{\log z}}{z}\sqrt{2^{1+depth_T(v)}n_v(\theta)}}{\sqrt{\frac{1}{2}n_v(\theta)\frac{\lfloor z^\alpha\rfloor}{z}}}=O\left(\frac{1}{z^{\frac{\alpha+\beta+1}{2}}}\right)+O\left(\frac{2^{\frac{depth_T(v)}{2}}\sqrt{\log z}}{z^{\frac{1+\alpha}{2}}}\right)$$

- $\sqrt{3\log z}\sqrt{2^{-depth_T(v)}}\sqrt{|\mathcal{I}|}> \frac{1}{2}2^{-depth_T(v)}|\mathcal{I}|$: It follows that

$$|\mathcal{I}|<12\;2^{depth_T(v)}\log z.$$

  Hence,

$$\frac{|\mu_1-\mu_2|}{\sigma_1}\le\frac{\frac{1}{z}+\frac{\sqrt{\log z}}{z}\sqrt{|\mathcal{I}|}}{\sqrt{\frac{1}{2}n_v(\theta)\frac{\lfloor z^\alpha\rfloor}{z}}}\le\frac{\frac{1}{z}+\frac{\sqrt{\log z}}{z}\sqrt{12\;2^{depth_T(v)}\log z}}{\sqrt{\frac{1}{2}n_v(\theta)\frac{\lfloor z^\alpha\rfloor}{z}}}=O\left(\frac{1}{z^{\frac{\alpha+\beta+1}{2}}}\right)+O\left(\frac{2^{\frac{depth_T(v)}{2}}\log z}{z^{\frac{\alpha+\beta+1}{2}}}\right)$$

Combining the above, it follows that

$$\begin{aligned} \big|\big|TP(\mu_1,\sigma_1^2)-TP(\mu_2,\sigma_2^2)\big|\big| &\le \frac{|\mu_1-\mu_2|}{\sigma_1}+\frac{|\sigma_1^2-\sigma_2^2|+1}{\sigma_1^2} \\ &\le O\left(\frac{1}{z^{\frac{\alpha+\beta+1}{2}}}\right)+O\left(\frac{2^{\frac{depth_T(v)}{2}}\sqrt{\log z}}{z^{\frac{1+\alpha}{2}}}\right)+O\left(\frac{2^{\frac{depth_T(v)}{2}}\log z}{z^{\frac{\alpha+\beta+1}{2}}}\right)+O(z^{-\alpha})+O(z^{-(\alpha+\beta-1)}) \\ &\le O\left(\frac{2^{\frac{depth_T(v)}{2}}\sqrt{\log z}}{z^{\frac{1+\alpha}{2}}}\right)+O\left(\frac{2^{\frac{depth_T(v)}{2}}\log z}{z^{\frac{\alpha+\beta+1}{2}}}\right)+O(z^{-\alpha})+O(z^{-(\alpha+\beta-1)}). \end{aligned}$$

Combining the above with (31) and (32) we get

$$\left\|\sum_{i\in\mathcal{N}_v(\theta)}Z_i-\sum_{i\in\mathcal{N}_v(\theta)}\widehat{Z}_i\right\|_{TV}\le O\left(\frac{2^{\frac{depth_T(v)}{2}}\sqrt{\log z}}{z^{\frac{1+\alpha}{2}}}\right)+O\left(\frac{2^{\frac{depth_T(v)}{2}}\log z}{z^{\frac{\alpha+\beta+1}{2}}}\right)+O(z^{-\alpha})+O(z^{-(\frac{\alpha+\beta-1}{2})}).$$

■

D Concentration of the Leaf Experiments

The following lemmas constitute the last piece of the puzzle and complete the proof of Lemma 3.5. They roughly state that, after the random walk in Stage 1 of the processes SAMPLING is performed, the experiments that will take place in Stage 2 of the processes SAMPLING are similar with high probability.

Proof of Lemma 3.10: Note that

$$\mu_v(\Phi) = \sum_{i \in \mathcal{I}} \Omega_i =: \Omega,$$

where $\{\Omega_i\}_i$ are independent random variables defined as

$$\Omega_i = \begin{cases} p_{i,v}(\ell_v^*), & \text{with probability } 2^{-depth_T(v)} \\ 0, & \text{with probability } 1 - 2^{-depth_T(v)}. \end{cases}$$

We apply the following version of Chernoff/Hoeffding bounds to the random variables $\Omega_i^{'} := z^{1-\alpha}\Omega_i \in [0,1]$.

Lemma D.1 (Chernoff/Hoeffding) *Let $Z_1, \ldots, Z_m$ be independent random variables with $Z_i \in [0,1]$, for all i. Then, if $Z = \sum_{i=1}^n Z_i$ and $\gamma \in (0,1)$,*

$$\Pr[|Z - \mathcal{E}[Z]| \geq \gamma \mathcal{E}[Z]] \leq 2\exp(-\gamma^2 \mathcal{E}[Z]/3).$$

Letting $\Omega^{'} = \sum_{i \in I} \Omega_i^{'}$ and applying the above lemma with $\gamma := \sqrt{\frac{1}{\mathcal{E}[\Omega']} \log z}$, it follows that

$$\Pr\left[\left|\Omega^{'} - \mathcal{E}[\Omega^{'}]\right| \geq \sqrt{\mathcal{E}[\Omega'] \log z}\right] \leq 2z^{-1/3},$$

which in turn implies

$$\Pr\left[|\Omega - \mathcal{E}[\Omega]| \geq z^{(\alpha-1)/2}\sqrt{\mathcal{E}[\Omega] \log z}\right] \leq 2z^{-1/3},$$

or, equivalently,

$$\Pr\left[|\mu_v(\Phi) - \mathcal{E}[\mu_v(\Phi)]| \geq z^{(\alpha-1)/2}\sqrt{\mathcal{E}[\mu_v(\Phi)] \log z}\right] \leq 2z^{-1/3}.$$

Similarly, it can be derived that

$$\Pr\left[\left|\widehat{\mu}_v(\hat{\Phi}) - \mathcal{E}[\widehat{\mu}_v(\hat{\Phi})]\right| \geq z^{(\alpha-1)/2}\sqrt{\mathcal{E}[\widehat{\mu}_v(\hat{\Phi})] \log z}\right] \leq 2z^{-1/3}.$$

Let us consider the joint probability space which makes $\Phi = \hat{\Phi}$ with probability 1; this space exists since as we observed above $G(\theta) = \widehat{G}(\theta), \forall \theta$. By a union bound for this space

$$\Pr\left[|\mu_v(\Phi) - \mathcal{E}[\mu_v(\Phi)]| \geq \frac{\sqrt{\log z}}{z^{(1-\alpha)/2}}\sqrt{\mathcal{E}[\mu_v(\Phi)]} \vee \left|\widehat{\mu}_v(\hat{\Phi}) - \mathcal{E}[\widehat{\mu}_v(\hat{\Phi})]\right| \geq \frac{\sqrt{\log z}}{z^{(1-\alpha)/2}}\sqrt{\mathcal{E}[\widehat{\mu}_v(\hat{\Phi})]}\right] \leq 4z^{-1/3}.$$

which implies

$$G\left(\theta : |\mu_v(\theta) - \mathcal{E}[\mu_v(\Phi)]| \leq \frac{\sqrt{\log z}}{z^{(1-\alpha)/2}}\sqrt{\mathcal{E}[\mu_v(\Phi)]} \wedge \left|\widehat{\mu}_v(\theta) - \mathcal{E}[\widehat{\mu}_v(\hat{\Phi})]\right| \leq \frac{\sqrt{\log z}}{z^{(1-\alpha)/2}}\sqrt{\mathcal{E}[\widehat{\mu}_v(\hat{\Phi})]}\right) \geq 1 - 4z^{-1/3}.$$

■

Proof of Lemma 3.11: Suppose that the random variables Φ and $\hat{\Phi}$ are coupled so that, with probability 1, $\Phi = \hat{\Phi}$. Then

$$\mu_v(\Phi) - \widehat{\mu}_v(\hat{\Phi}) = \sum_{i \in \mathcal{I}} \Omega_i =: \Omega,$$

where $\{\Omega_i\}_i$ are independent random variables defined as

$$\Omega_i = \begin{cases} p_{i,v}(\ell_v^*) - \widehat{p}_{i,v}(\ell_v^*), & \text{with probability } 2^{-depth_T(v)} \\ 0, & \text{with probability } 1 - 2^{-depth_T(v)}. \end{cases}$$

We apply Hoeffding's inequality to the random variables Ω_i.

Lemma D.2 (Hoeffding's Inequality) *Let $X_1, \ldots, X_n$ be independent random variables. Assume that, for all i, $\Pr[X_i \in [a_i, b_i]] = 1$. Then, for $t > 0$:*

$$\Pr\left[\sum_i X_i - \mathcal{E}\left[\sum_i X_i\right] \geq t\right] \leq \exp\left(-\frac{2t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right).$$

Applying the above lemma we get

$$\Pr\left[|\Omega - \mathcal{E}\left[\Omega\right]| \geq t\right] \leq 2\exp\left(-\frac{2t^2}{|\mathcal{I}|\frac{4}{z^2}}\right),$$

since, for all $i \in \mathcal{I}$, $|p_{i,v}(\ell_v^*) - \widehat{p}_{i,v}(\ell_v^*)| \leq \frac{1}{z}$. Setting $t = \sqrt{\log z}\sqrt{|\mathcal{I}|}\frac{1}{z}$ we get

$$\Pr\left[|\Omega - \mathcal{E}\left[\Omega\right]| \geq \sqrt{\log z}\sqrt{|\mathcal{I}|}\frac{1}{z}\right] \leq 2\frac{1}{z^{1/2}}.$$

Note that

$$|\mathcal{E}[\Omega]| = |\sum_{i \in \mathcal{I}} \mathcal{E}[\Omega_i]| = |2^{-depth_T(v)} \sum_{i \in \mathcal{I}} (p_{i,v}(\ell_v^*) - \widehat{p}_{i,v}(\ell_v^*))| \leq \frac{1}{z}.$$

It follows from the above that

$$\Pr\left[|\Omega| \leq \frac{1}{z} + \sqrt{\log z}\sqrt{|\mathcal{I}|}\frac{1}{z}\right] \geq 1 - 2\frac{1}{z^{1/2}},$$

which gives immediately that

$$G\left(\theta : |\mu_v(\theta) - \widehat{\mu}_v(\theta)| \leq \frac{1}{z} + \frac{\sqrt{\log z}}{z}\sqrt{|\mathcal{I}|}\right) \geq 1 - \frac{2}{z^{1/2}}.$$

Moreover, an easy application of Lemma D.1 gives

$$G\left(\theta : |n_v(\theta) - 2^{-depth_T(v)}|\mathcal{I}|| \leq \sqrt{3\log z}\sqrt{2^{-depth_T(v)}|\mathcal{I}|}\right) \geq 1 - \frac{2}{z}. \tag{33}$$

Indeed, let $T_i = 1_{\Phi_i = v}$. Then $n_v(\Phi) = \sum_{i \in \mathcal{I}} T_i$ and $\mathcal{E}[\sum_{i \in \mathcal{I}} T_i] = 2^{-depth_T(v)}|\mathcal{I}|$. Applying Lemma D.1 with $\gamma = \sqrt{\frac{3\log z}{2^{-depth_T(v)}|\mathcal{I}|}}$ we get

$$\Pr\left[\left|\sum_{i \in \mathcal{I}} T_i - \mathcal{E}\left[\sum_{i \in \mathcal{I}} T_i\right]\right| \geq \sqrt{3\log z}\sqrt{2^{-depth_T(v)}|\mathcal{I}|}\right] \leq \frac{2}{z},$$

which implies

$$\Pr\left[|n_v(\Phi) - 2^{-depth_T(v)}|\mathcal{I}|| \leq \sqrt{3\log z}\sqrt{2^{-depth_T(v)}|\mathcal{I}|}\right] \geq 1 - \frac{2}{z}.$$

■